DOCUMENT RESUME

ED 063 267 24 SP 005 739

AUTHOR Atkinson, Richard C.
TITLE Methods for Maximizing the Learning Process: A Theoretical and Experimental Analysis.
BUREAU NO BR-9-0401
PUB DATE Aug 71
GRANT OEG-0-9-140401-4147(010)
NOTE 48p.

EDRS PRICE MF-$0.65 HC-$3.29
DESCRIPTORS *Instructional Design; *Instructional Improvement; *Instructional Innovation; *Learning Processes; *Learning Theories; Mathematical Models

ABSTRACT

This research deals with optimizing the instructional process. The approach adopted was to limit consideration to simple learning tasks for which adequate mathematical models could be developed. Optimal or suitable suboptimal instructional strategies were developed for the models. The basic idea was to solve for strategies that either maximize the amount learned in a fixed time period or minimize the time necessary to attain a prescribed level of performance. Once strategies were formulated, experiments were carried out to evaluate their relative efficiency. The program of work invclved a mathematical analysis of optimization problems relating to the learning process and also representing a fairly unique method for testing learning theories. In this sense, the project was an attempt to bridge the gap between the psychologist's laboratory experiments in learning theory and the practical problems of divising efficient instructional strategies in the classroom. The optimization strategies developed and tested during the course of the project were fairly restrictive in character and applicable primarily to simple tasks such as those found in initial reading, language arts, and elementary mathematics. An 18-item bibliography and appendixes are included. (Author/MJM)

FILMED FROM BEST AVAILABLE COPY

PA 24
9-0401

ED 063267

Methods for Maximizing the Learning Process:

A Theoretical and Experimental Analysis

FINAL REPORT

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
OFFICE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL OFFICE OF EDUCATION POSITION OR POLICY.

Project No. 9-0401
Grant No. OEG-0-9-140401-4147(010)

Richard C. Atkinson

August 1971

The research reported herein was performed pursuant to a grant with the Office of Education, U.S. Department of Health, Education, and Welfare. Contractors undertaking such projects under Government sponsorship are encouraged to express freely their professional judgment in the conduct of the project. Points of view or opinions stated do not, therefore, necessarily represent official Office of Education position or policy.

U.S. DEPARTMENT OF HEALTH, EDUCATION, AND WELFARE
Office of Education
National Center for Educational
Research and Development
(Regional Research Program)

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

SP 005 739

CONTENTS

	Page
Summary	1
Introduction	1
A Decision-Theoretic Analysis of Instruction	2
Criteria for a Theory of Instruction	3
Optimizing Instruction in Initial Reading	4
Optimal Sequencing Procedures	8
An Evaluation of the All-Or-None Strategy	14
Conclusions and Recommendations	20
References	24
Appendix I	26

SUMMARY

The research reported here deals with the topic of optimizing the instructional process. The problem can be investigated in many ways, but the approach adopted was to limit consideration primarily to simple learning tasks for which adequate mathematical models could be developed and shown to be reasonably accurate.

For these models, we derived optimal or suitable suboptimal instructional strategies. The basic idea was to solve for strategies that either maximize the amount learned in a fixed time period or minimize the time necessary to attain a prescribed level of performance. Once such strategies had been formulated, experiments were carried out to evaluate their relative efficiency. To the extent that particular strategies proved effective, they were incorporated into computer-based instructional programs in initial reading currently in operation at Stanford University.

The program of work involved a mathematical analysis of optimization problems related to the learning process, and also represented a fairly unique method for testing theories of learning. In this sense the project was an attempt to bridge the gap between the psychologist's laboratory experiments in learning theory and the practical problems of devising efficient instructional strategies in the classroom. The optimization strategies developed and tested during the course of this project were fairly restrictive in character and applicable primarily to simple tasks such as those found in initial reading, language arts, and elementary mathematics. On the other hand, it is our hope that mathematically precise models for optimizing learning in these simple tasks may in time provide guidelines for a theory of instruction that is mathematically precise and yet has wide applicability.

INTRODUCTION

The term "theory of instruction" has been in widespread use for over a decade and during that time has acquired a fairly specific meaning. By consensus it denotes a body of theory concerned with optimizing the learning process; stated otherwise, the goal of a theory of instruction is to prescribe the most effective methods for acquiring new information, whether in the form of higher-order concepts or rote facts. Although usage of the term is widespread, there is no agreement on the requirements for a theory of instruction. The literature provides an array of examples ranging from speculative accounts of how children should be taught in the classroom to formal mathematical models specifying precise branching procedures in computer-controlled instruction. Such diversity is healthy; to focus on only one approach would not be productive in the long run. I prefer to use the term "theory of instruction" to encompass both experimental and theoretical research, with the theoretical work ranging from general speculative accounts to specific quantitative models.

The task of going from a description of the learning process to a prescription for optimizing learning must be clearly distinguished from the task of finding the appropriate theoretical description in the first place. However, there is a danger that preoccupation with finding prescriptions for instruction may cause us to overlook the critical interplay between the

two interprises. Developments in control theory and statistical decision theory provide potentially powerful methods for discovering optimal decision-making strategies in a wide variety of contexts. In order to use these tools it is necessary to have a reasonable model of the process to be optimized. Some learning processes can already be described with the required degree of accuracy. This report will examine an approach to the psychology of instruction which is appropriate when the learning is governed by such a process.

A DECISION-THEORETIC ANALYSIS OF INSTRUCTION

The derivation of an optimal strategy requires that the instructional problem be stated in a form amenable to a decision-theoretic analysis. Analyses based on decision theory vary somewhat from field to field, but the same formal elements can be found in most of them. As a starting point it will be useful to identify these elements in a general way, and then relate them to an instructional situation. They are as follows:

1. The possible states of nature.
2. The actions that the decision-maker can take to transform the state of nature.
3. The transformation of the state of nature that results from each action.
4. The cost of each action.
5. The return resulting from each state of nature.

In the context of instruction, these elements divide naturally into three groups. Elements 1 and 3 are concerned with a description of the learning process; elements 4 and 5 specify the cost-benefit dimensions of the problem; and element 2 requires that the instructional actions from which the decision maker is free to chose be precisely specified.

For the decision problems that arise in instruction, elements 1 and 3 require that a model of the learning process exist. It is usually natural to identify the states of nature with the learning states of the student. Specifying the transformation of the states of nature caused by the actions of the decision-maker is tantamount to constructing a model of learning for the situation under consideration. The learning model will be probabilistic to the extent that the state of learning is imperfectly observable or the transformation of the state of learning that a given instructional action will cause is not completely predictable.

The specification of costs and returns in an instructional situation (elements 4 and 5) tends to be straightforward when examined on a short-term basis, but virtually intractable over the long-term. For the short-term one can assign costs and returns for the mastery of, say, certain basic reading skills, but sophisticated determinations for the long-term value of these skills to the individual and society are difficult to make. There is an important role for detailed economic analyses of the long-term impact of education, but such studies deal with issues at a more global level than we shall consider here. The present analysis will be limited to those costs and returns directly related to a specific instructional task.

Element 2 is critical in determining the effectiveness of a decision-theory analysis; the nature of this element can be indicated by an example. Suppose we want to design a supplementary set of exercises for an initial reading program that involve both sight-word identification and phonics. Let us assume that two exercise formats have been developed, one for training on sight words, the other for phonics. Given these formats, there are many ways to design an overall program. A variety of optimization problems can be generated by fixing some features of the curriculum and leaving others to be determined in a theoretically optimal manner. For example, it may be desirable to determine how the time available for instruction should be divided between phonics and sight word recognition, with all other features of the curriculum fixed. A more complicated question would be to determine the optimal ordering of the two types of exercises in addition to the optimal allocation of time. It would be easy to continue generating different optimization problems in this manner. The main point is that varying the set of actions from which the decision-maker is free to choose changes the decision problem, even though the other elements remain the same.

Once these five elements have been specified, the next task is to derive the optimal strategy for the learning model that best describes the situation. If more than one learning model seems reasonable a priori, then competing candidates for the optimal strategy can be deduced. When these tasks have been accomplished, an experiment can be designed to determine which strategy is best. There are several possible directions in which to proceed after the initial comparison of strategies, depending on the results of the experiment. If none of the supposedly optimal strategies produces satisfactory results, then further experimental analysis of the assumptions of the underlying learning models is indicated. New issues may arise even if one of the procedures is successful. In one of the experiments that we shall report, the successful strategy produces an unusually high error rate during learning, which is contrary to a widely accepted principle of programmed instruction (Skinner, 1968). When anomalies such as this occur, they suggest new lines of experimental inquiry, and often require a reformulation of the learning model.

CRITERIA FOR A THEORY OF INSTRUCTION

Our discussion to this point can be summarized by listing four criteria that must be satisfied prior to the derivation of an optimal instructional strategy.

1. A model of the learning process.
2. Specification of admissible instructional actions.
3. Specification of instructional objectives.
4. A measurement scale that permits costs to be assigned to each of the instructional actions and payoffs to the achievement of instructional objectives.

If these four elements can be given a precise interpretation then it is generally possible to derive an optimal instructional policy. The solution for an optimal policy is not guaranteed, but in recent years some powerful tools have been developed for discovering optimal or near optimal procedures if they exist.

The four criteria listed above, taken in conjunction with methods for deriving optimal strategies, define either a model of instruction or a theory of instruction. Whether the term theory or model is used depends on the generality of the applications that can be made. Much of my own work has been concerned with the development of specific models for specific instructional tasks; hopefully, the collection of such models will provide the groundwork for a general theory of instruction.

In terms of the criteria listed above, it is clear that a model or theory of instruction is in fact a special case of what has come to be known in the mathematical and engineering literature as optimal control theory or, more simply, control theory (Kalman, Falb, & Arbib, 1969). The development of control theory has progressed at a rapid rate both in the United States and abroad, but most of the applications involve engineering or economic systems of one type or another. Precisely the same problems are posed in the area of instruction except that the system to be controlled is the human learner, rather than a machine or group of industries. To the extent that the above four elements can be formulated explicitly, methods of control theory can be used in deriving optimal instructional strategies.

In the experiments that we shall report, two basic types of strategies are examined. One is a response-insensitive strategy and the other a response-sensitive strategy. A response-insensitive strategy orders the instructional materials without taking into account the student's responses (except possibly to provide corrective feedback) as he progresses through the curriculum. In contrast, a response-sensitive strategy makes use of the student's response history in its stage-by-stage decisions regarding which curriculum materials to present next. Response-insensitive strategies are completely specified in advance and consequently do not require a system capable of branching during an instructional session. Response-sensitive strategies are more complex, but have the greatest promise for producing significant gains for they must be at least as good, if not better, than the comparable response-insensitive strategy.

OPTIMIZING INSTRUCTION IN INITIAL READING

The first study to be described here is based on work concerned with the development of a computer-assisted instruction (CAI) program for teaching reading in the primary grades (Atkinson & Fletcher, 1972). The program provides individualized instruction in reading and is used as a supplement to normal classroom teaching; a given student may spend anywhere from zero to 30 minutes per day at a CAI terminal. For present purposes only one set of results will be considered, where the dependent measure is performance on a standardized reading achievement test administered at the end of the first grade. Using our data a statistical model can be formulated that predicts test performance as a function of the amount of time the student spends on the CAI system. Specifically, let $P_i(t)$ be student i's performance on a reading test administered at the end of first grade, given that he spends time t on the CAI system during the school year. Then within certain limits the following equation holds:

$$P_i(t) = \alpha_i - \beta_i \exp(-\gamma_i t) \tag{1}$$

Depending on a student's particular parameter values, the more time spent on the CAI program the higher the level of achievement at the end of the year. The parameters α, β, and γ, characterize a given student and vary from one student to the next; α and $(\alpha-\beta)$ are measures of the student's maximal and minimal levels of achievement respectively, and γ is a rate of progress measure. These parameters can be estimated from a student's response record obtained during his first hour of CAI. Stated otherwise, data from the first hour of CAI can be used to estimate the parameters α, β, and γ for a given student, and then the above equation enables us to predict end-of-year performance as a function of the CAI time allocated to that student.

The optimization problem that arises in this situation is as follows: Let us suppose that a school has budgeted a fixed amount of time T on the CAI system for the school year and must decide how to allocate the time among a class of n first-grade students. Assume, further, that all students have had a preliminary run on the CAI system so that estimates of the parameters α, β, and γ have been obtained for each student.

Let t_i be the time allocated to student i. Then the goal is to select a vector $(t_1, t_2,\ldots,t_n)$ that optimizes learning. To do this let us check our four criteria for deriving an optimal strategy.

The first criterion is that we have a model of the learning process. The prediction equation for $P_i(t)$ does not offer a very complete account of learning; however, for purposes of this problem the equation suffices as a model of the learning process, giving all of the information that is required. This is an important point to keep in mind: the nature of the specific optimization problem determines the level of complexity that must be represented in the learning model. For some problems the model must provide a relatively complete account of learning in order to derive an optimal strategy, but for other problems a simple descriptive equation of the sort presented above will suffice.

The second criterion requires that the set of admissible instructional actions be specified. For the present case the potential actions are simply all possible vectors $(t_1, t_2,\ldots,t_n)$ such that the t_i's are non-negative and sum to T. The only freedom we have as decision makers in this situation is in the allocation of CAI time to individual students.

The third criterion requires that the instructional objective be specified. There are several objectives that we could choose in this situation. Let us consider four possibilities:

(a) Maximize the mean value of P over the class of students.
(b) Minimize the variance of P over the class of students.
(c) Maximize the number of students who score at grade level at the end of the first year.
(d) Maximize the mean value of P satisfying the constraint that the resulting variance of P is less than or equal to the variance that would have been obtained if no CAI was administered.

Objective (a) maximizes the gain for the class as a whole: (b) aims to reduce differences among students by making the class as homogeneous as

possible; (c) is concerned specifically with those students that fall behind grade level; (d) attempts to maximize performance of the whole class but insures that differences among students are not amplified by CAI. Other instructional objectives can be listed, but these are the ones that seemed most relevent. For expository purposes, let us select (a) as the instructional objective.

The fourth criterion requires that costs be assigned to each of the instructional actions and that payoffs be specified for the instructional objectives. In the present case we assume that the cost of CAI does not depend on how time is allocated among students and that the measurement of payoff is directly proportional to the students' achieved value of P.

In terms of our four criteria, the problem of deriving an optimal instructional strategy reduces to maximizing the function

$$(2) \qquad \phi(t_1, t_2, \ldots, t_n) = \frac{1}{n} \sum_{i=1}^{n} P_i(t_i)$$

$$= \frac{1}{n} \sum_{i=1}^{n} \alpha_i + \beta_i \exp(-\gamma_i t_i)$$

subject to the constraint that

$$(3) \qquad \sum_{i=1}^{n} t_i = T$$

and

$$t_i \geq 0.$$

This maximization can be done by using the method of dynamic programming (Bellman, 1961). In order to illustrate the approach, computations were made for a first-grade class where the parameters α, β, and γ had been estimated for each student. Employing these estimates, computations were carried out to determine the time allocations that maximized the above equation. For the optimal policy the predicted mean performance level of the class, $\bar{P}$, was 15% higher than a policy that allocated time equally to students (i.e., a policy where $t_i = t_j$ for all i and j). This gain represents a substantial improvement; the drawback is that the variance of the P scores is roughly 15% greater than for the equal-time policy. This means that if we are interested primarily in raising the class average, we must let the rapid learners move ahead and progress far beyond the slow learners.

Although a time allocation that complies with objective (a) did increase overall class performance, the correlated increase in variance leads us to believe that other objectives might be more beneficial. For comparison, time allocations also were computed for objectives (b), (c), and (d). Figure 1 presents the predicted gain in $\bar{P}$ as a percentage of $\bar{P}$ for the equal-time

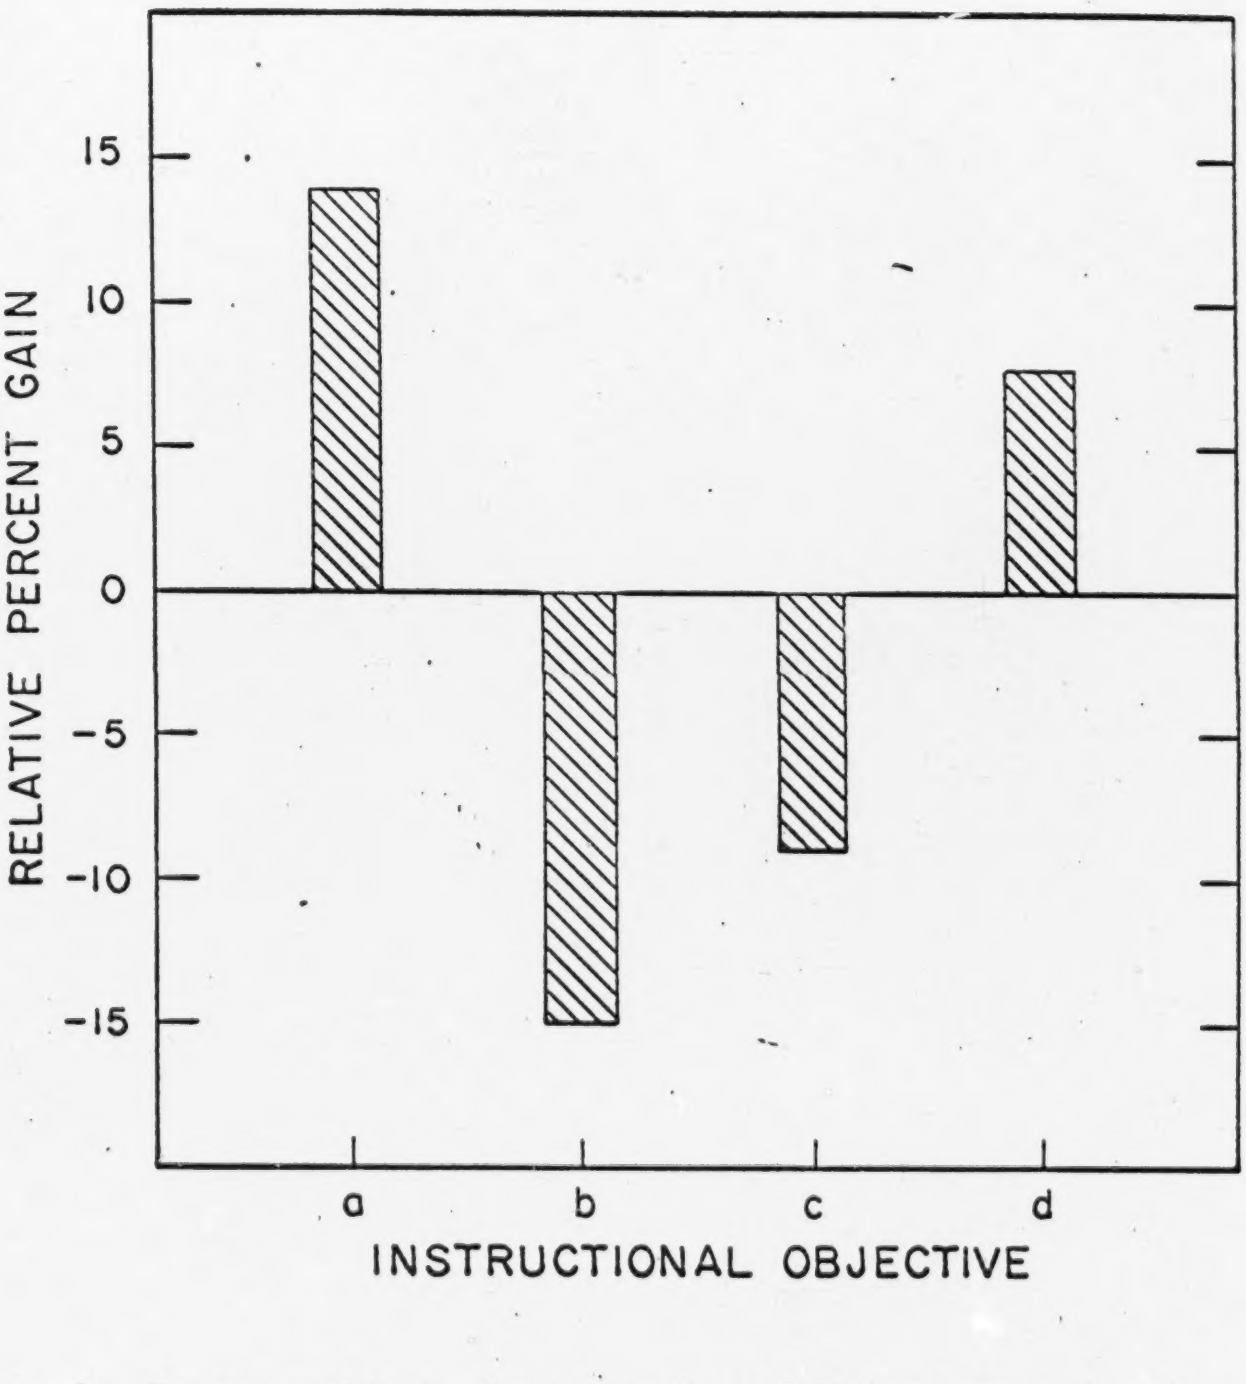

Figure 1: Percent gains in the mean value of P when compared with an equal-time policy for four policies each based on a different instructional objective.

policy. Objectives (b) and (c) yield negative gains and so they should since their goal is to reduce variability, which is accomplished by holding back on the rapid learners and giving a lot of attention to the slower ones. The reduction in variability for these two objectives, when compared with the equal-time policy, is 12% and 10%, respectively. Objective (d), which attempts to strike a balance between objective (a) on the one hand and objectives (b) and (c) on the other, yields an 8% increase in $\overline{P}$ and yet reduces variability by 6%.

In view of these computations, objective (d) seems to be preferred; it offers a substantial increase in mean performance while maintaining a low level of variability. As yet, we have not implemented this policy, so only theoretical results can be reported. Nevertheless, these examples yield differences that illustrate the usefulness of this type of analysis. They make it clear that the selection of an instructional objective should not be done in isolation, but should involve a comparative analysis of several alternatives taking into account more than one dimension of performance. For example, even if the principal goal is to maximize $\overline{P}$, it would be inappropriate in most educational situations to select a given objective over some other if it yielded only a small average gain while variability mushroomed.

OPTIMAL SEQUENCING PROCEDURES

One application of computer-assisted instruction (CAI) which has proved to be very effective in the primary grades involves a regular program of practice and review specifically designed to complement the efforts of the classroom teacher (Atkinson, 1969). Some of the curriculum materials in such programs take the form of lists of instructional units or items. The objective of the CAI programs is to teach students the correct response to each item in a given list. Typically, a sublist of items is presented each day in one or more fixed exercise formats. The optimization problem that arises concerns the selection of items for presentation on a given day.

The Stanford Reading Project is an example of such a program in initial reading instruction (Atkinson, Fletcher, Chetin, & STauffer, 1971). The vocabularies of several of the commonly used basal readers were compiled into one dictionary and a variety of exercises using these words were developed to teach reading skills. These exercises were designed principally to strengthen the student's decoding skills, with special emphasis on letter identification, sight-word recognition, phonics, spelling patterns, and word comprehension. The details of the teaching procedure vary from one exercise to another, but most include a sequence in which a curriculum item is presented, eliciting a response from the student, followed by a short period for studying the correct response. For example, one exercise in sight-word recognition has the following format:

Teletype Display	Audio Message
NUT MEN RED	Type red

Three words are printed on the teletype, followed by an audio presentation of one of the words. Control is then turned over to the student; if he types the

correct word a reinforcing message is given and the computer program then proceeds to the next presentation. If the student responds incorrectly or exceeds the time, the teletype prints the correct word simultaneously with its audio presentation and then moves to the next presentation. Under an early version of the program, items were presented in a predetermined sublists, with an exercise continuing on a sublist until a specified criterion has been met.

Strategies can be found that will improve on the fixed order of presentation. Two studies to be described below are concerned with the development of such strategies. One study examines alternative presentation strategies for teaching spelling words to elementary school children, and the other examines strategies for teaching Swahili vocabulary items to college-level students. The optimization problems in both studies were essentially the same. A list of N items is to be learned, and a fixed number of days, D, are allocated for its study. On each day a sublist of items is presented for test and study. The sublist always involves M items and each is presented only once for test followed by a study period. The total set of N items is extremely large with regard to any sublist of M items. Once the experimenter has specified a sublist for a given day its order of presentation is random. After the D days of study are completed, a post-test is given over all items. The parameters N, D and M are fixed, and so is the instructional format on each day. Within these constraints the problem is to maximize performance on a post-test by an appropriate selection of sublists from day to day. The strategy for selecting sublists from day to day is dynamic (or response sensitive, using the terminology of Groen and Atkinson, 1966) to the extent that it depends upon the student's prior history of performance.

Three Models of the Learning Process

Two extremely simple learning models will be considered first. Then a third model which combines features of the first two will be described.

In the first model, the state of the learner with respect to each item is completely determined by the number of times the item has been studied. At the start of the experiment an item has some initial probability of error; each time the item is presented its error probability is reduced by a factor α, which is less than one. Stated as a difference equation, the probability of an error on the $n+1^{st}$ presentation of an item is related to its probability on the n^{th} presentation as follows:

$$q_{n+1} = \alpha q_n. \tag{4}$$

Note that the error probability for a given item depends on the number of times it has been reduced by the factor α; that is, the number of times it has been presented. Learning is the gradual reduction in the probability of error by repeated presentations of items. This model is sometimes called the linear model because the equation describing change in response probability is linear.

In the second model, mastery of an item is not at all gradual. At any point in time a student is in one of two states with respect to each item: the learned state or the unlearned state. If an item in the learned state is presented, the correct response is always given; if an item is in the unlearned

state, an incorrect response is given unless the student makes a correct response by guessing. When an unlearned item is presented, it may move into the learned state with probability c. Stated as a difference equation,

$$q_{n+1} = \begin{cases} q_n & \text{, with probability } 1-c \\ 0 & \text{, with probability } c \text{ .} \end{cases} \tag{5}$$

Once an item is learned, it remains in the learned state throughout the course of instruction. Some items are learned the first time they are presented, others may be presented several times before they are finally learned. Therefore, the list as a whole is learned gradually. But for any particular item, the transition from the unlearned to the learned state occurs on a single trial. The model is sometimes called the all-or-none-model because of this characterization of the possible states of learning.

The third model to be considered is the random-trial increments (RTI) model and represents a compromise between the linear and all-or-none model (Norman, 1964). For this model

$$q_{n+1} = \begin{cases} q_n & \text{, with probability } 1-c \\ \alpha q_n & \text{, with probability } c \text{ .} \end{cases} \tag{6}$$

If $c = 1$, the RTI model reduces to the linear model; if $\alpha = 0$, it reduces to the all-or-none model. However, for $c < 1$ and $\alpha > 0$, the RTI model generates predictions that are quite distinct from both the linear and the all-or-none models.

For all three models the probability of an error on the first trial is a parameter that may need to be estimated in certain situations; to emphasize this point the initial error probability will be written as q' henceforth. It should be noted that the all-or-none model and the RTI model are response sensitive in that the learner's particular history of correct and incorrect responses makes a difference in predicting performance on the next presentation of an item. In contrast, the linear model is response insensitive; its prediction depends only on the number of prior presentations and is not improved by a knowledge of the learner's response history.

The Cost/Benefit Structure

At the present level of analysis, it will expedite matters if some assumptions are made to simplify the appraisal of costs and benefits associated with various strategies. It is tacitly assumed that the subject matter being taught is sufficiently beneficial to justify allocating a fixed amount of time to it for instruction. Since the exercise formats and the time allocated to instruction are the same for all strategies, it is reasonable to assume that the costs of instruction are the same for all strategies as well. If the

costs of instruction are equal for all strategies, then for purposes of comparison they may be ignored and attention focused on the comparative benefits of the various strategies. This is an important simplification because it affects the degree of precision necessary in the assessment of costs and benefits. If both costs and benefits are significantly variable in a problem, then it is essential that both quantities be estimated accurately. This is often difficult to do. When one of these quantities can be ignored, it suffices if the other can be assessed accurately enough to order the possible outcomes. This is usually fairly easy to accomplish. In the present problem, for example, it is reasonable to consider all the items equally beneficial. This implies that benefits depend only on the overall probability of a correct response, not on the particular items known. It turns out that this specification of cost and benefit is sufficient for the learning models to determine optimal strategies.

The above cost/benefit assumptions permit us to concentrate on the main concern of this paper, the derivation of the educational implications of learning models. Also, they are approximately valid in many instructional contexts. Nevertheless, it must be recognized that in the majority of cases these assumptions will not be satisfied. For instance, the assumption that the alternative strategies cost the same to implement usually does not hold. It only holds as a first approximation in the case being considered here. In the present formulation of the problem, a fixed amount of time is allocated for study and the problem is to maximize learning, subject to this time constraint. An alternative formulation which is more appropriate in some situations fixes a minimum criterion level for learning. In this formulation, the problem is to find a strategy for achieving this criterion level of performance in the shortest time. As a rule, both costs and benefits must be weighed in the analysis, and frequently subtopics within a curriculum vary significantly in their importance. Sometimes there is a choice among several exercise formats. In certain cases, whether or not a certain topic should be taught at all is the critical question. Smallwood (1971) has treated a problem similar to the one considered in this paper in a way that includes some of these factors in the structure of costs and benefits.

Deducing Strategies from the Learning Models

Optimal strategies can be deduced for the linear and all-or-none models under the assumption that all items have ghe same learning parameters and initial error probabilities. The situation is more complicated in the case of the RTI model. An approximation to the optimal strategy for the RTI case will be discussed later; in this case the strategy explicitly allows for individual differences in parameter values.

For the linear model, if an item has been presented n times, the probability of an error on the next presentation of the item is $\alpha^{n-1}q'$; when the item is presented, the error probability is reduced to $\alpha^{n}q'$. The size of the reduction is thus $\alpha^{n-1}(1-\alpha)q'$. Observe that the size of the decrement in error probability gets smaller with each presentation of the item. This observation can be used to deduce that the following procedure is optimal.

> On a given day, form the sublist of M items by selecting those items that have received the fewest presentations up to that point. If more than M items satisfy this criterion, then select items at random from the set satisfying the criterion.

Upon examination, this strategy is seen to be equivalent to the standard cyclic presentation procedure commonly employed in experiments on paired-associate learning. It amounts to presenting all items once, randomly reordering them, presenting them again and repeating the procedure until the number of days allocated to instruction have been exhausted.

According to the all-or-none model, once an item has been learned there is no further reason to present it. Since all unlearned items are equally likely to be learned if presented, it is intuitively reasonable that the optimal presentation strategy selects the item least likely to be in the learned state for presentation. In order to discover a good index of the likelihood of being in the learned state, consider a student's response protocol for a single item. If the last response was incorrect, the item was certainly in the unlearned state at that time, although it may then have been learned during the study period that immediately followed. If the last response was correct, then it is more likely that the item is now in the learned state. In general, the more correct responses there are in the protocol since the last error on the item, the most likely it is that the item is in the learned state.

The preceding observations provide a heuristic justification for an algorithm which Karush and Dear (1966) have proved is in fact the optimal strategy for the all-or-none model. The optimal strategy requires that for each student a bank of counters be set up, one for each word in the list. To start, M different items are presented each day until each item has been presented once and a 0 has been entered in its counter. On all subsequent days the strategy requires that we conform to the following two rules:

1. Whenever an item is presented, increase its counter by 1 if the subject's response is correct, but reset it to 0 if the response is incorrect.

2. Present the M items whose counters are lowest among all items. If more than M items are eligible, then select randomly as many items as are needed to complete the sublist of size M from those having the same highest counter reading, having selected all items with lower counter values.

For example, suppose 6 items are presented each day and after a given day a certain student has 4 items whose counters are 0, 4 whose counters are 1, and higher values for the rest of the counters. His study list would consist of the 4 items whose counters are 0, and 2 items selected at random from the 4 whose counters are 1.

It has been possible to find relatively simple optimal strategies for the linear and all-or-none models. It is noteworthy that neither strategy depends on the values of the parameters of the respective models (i.e., on α, c, or q'). Another exceptional feature of these two models is that it is possible to condense a student's response protocol to one index per item without losing any information relevant to presentation decisions. Such condensations of response protocols are referred to as sufficient histories (Groen & Atkinson, 1966). Roughly speaking, an index summarizing the information in a student's response protocol is a sufficient history if any additional information from the protocol would be redundant in the determination of the student's state of learning. The concept is analogous to a sufficient statistic. If one takes a sample of observations from a population with an underlying normal distribution and wishes to estimate the population mean, the sample mean is a sufficient statistic. Other statistics that can be calculated (such as the median, the range, and the standard deviation) cannot be used to improve on the sample mean as an estimate of the population mean, though they may be useful in assessing the precision of the estimate. In statistics, whether or not data can be summarized by a few simple sufficient statistics is determined by the nature of the underlying distribution. For educational applications, whether or not a given instructional process can be adequately monitored by a simple sufficient history is determined by the model representing the underlying learning process.

The random-trial increments model appears to be an example of a process for which the information in the subject's response protocol cannot be condensed into a simple sufficient history. It is also a model for which the optimal strategy depends on the values of the model parameters. Consequently, it is not possible to state a simple algorithm for the optimal presentation strategy for this model. Suffice it to say that there is an easily computable formula for determining which item has the best expected immediate gain, if presented. The strategy that presents this item should be a reasonable approximation to the optimal strategy (Calfee, 1970). More will be said later regarding the problem of parameter estimation and some of its ramifications.

If the three models under consideration are to be ranked on the basis of their ability to account for data from laboratory experiments employing the standard presentation procedure, the order of preference is clear. The all-or-none model provides a better account of the data than the linear model, and the random-trial increments model is better than either of them (Atkinson & Crothers, 1964). This does not necessarily imply, however, that the optimization strategies derived from these models will receive the same ranking. The standard cyclic presentation procedure used in most learning experiments may mask certain deficiencies in the all-or-none or RTI models which would manifest themselves when the optimal presentation strategy specified by one or the other of these models was employed.

AN EVALUATION OF THE ALL-OR-NONE STRATEGY

Lorton (1971), in a Ph.D. Thesis conducted under the auspices of the present grant, compared the all-or-none strategy with the standard procedure in an experiment in computer-assisted spelling instruction with elementary school children. The former strategy is optimal if the learning process is indeed all-or-none, whereas the latter is optimal if the process is linear. The experiment was one phase of the Stanford Reading Project using computer facilities at Stanford University linked via telephone lines to student terminals in the schools.

Individual lists of 48 words were compiled in an extensive pretest program to guarantee that each student would be studying words of approximately equal difficulty which he did not already know how to spell. A within-subjects design was used in an effort to make the comparison of strategies as sensitive as possible. Each student's individualized list of 48 words was used to form two comparable lists of 24 words, one to be taught using the all-or-none strategy and the other using the standard procedure. Each day a student was given training on 16 words, 8 from the list for standard presentation and 8 from the list for presentation according to the all-or-none strategy. There were 24 training sessions followed by three days for testing all the words; approximately two weeks later three more days were spent on a delayed retention test. Using this procedure, all words in the standard presentation list received exactly one presentation in successive 3-day blocks during training. Words in the list presented according to the all-or-none algorithm received from 0 to 3 presentations in successive 3-day blocks during training, with one presentation being the average. A flow chart of the daily routine is given in Figure 2.

The results of the experiment are summarized in Figure 3. The proportions of correct responses are plotted for successive 3-day blocks during training, followed by the first overall test and then the two-week delayed test. Note that during training the proportion correct is always lower for the all-or-none procedure than for the standard procedure, but on both the final test and the retention test the proportion correct is greater for the all-or-none strategy. Analysis of variance tests verified that these results are statistically significant. The advantage of approximately ten percentage points on the post-tests for the all-or-none procedure is of practical significance as well.

The observed pattern of results is exactly what would be predicted if the all-or-none model does indeed describe the learning process. As was shown earlier, final test performance should be better when the all-or-none optimization strategy is adopted as opposed to the standard procedure. Also the greater proportion of error for this strategy during training is to be expected. The all-or-none strategy presents the items least likely to be in the learned state, so it is natural that more errors would be made during training.

A TEST OF A PARAMETER-DEPENDENT STRATEGY

As noted earlier, the strategy derived for the all-or-none model in the case of homogeneous items does not depend on the actual values of the model parameters. In many situations either the assumptions of the all-or-none model or the assumption of homogeneous items or both are seriously violated,

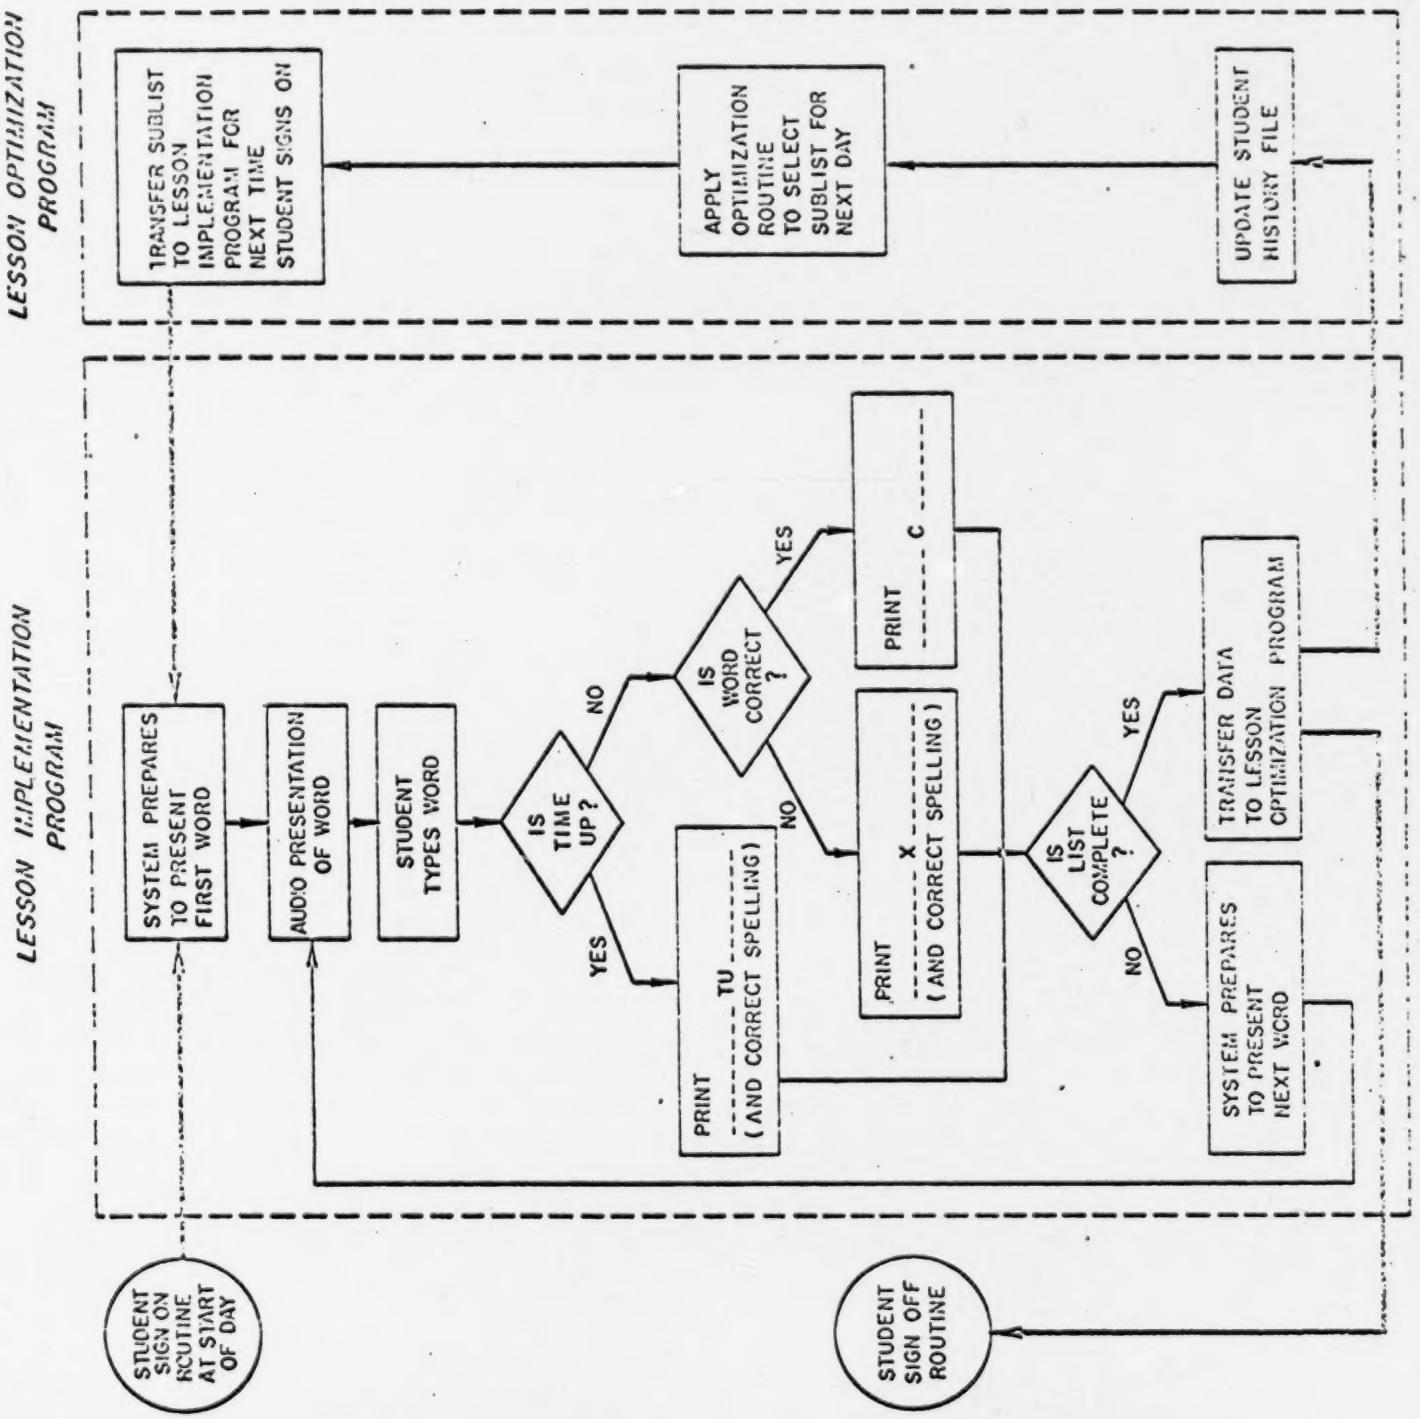
LESSON OPTIMIZATION PROGRAM
TRANSFER SUBLIST TO LESSON IMPLEMENTATION PROGRAM FOR NEXT TIME STUDENT SIGNS ON
APPLY OPTIMIZATION ROUTINE TO SELECT SUBLIST FOR NEXT DAY
UPDATE STUDENT HISTORY FILE
LESSON IMPLEMENTATION PROGRAM
SYSTEM PREPARES TO PRESENT FIRST WORD
AUDIO PRESENTATION OF WORD
STUDENT TYPES WORD
IS TIME UP?
NO
YES
IS WORD CORRECT ?
YES
NO
PRINT ---------- C ----------
PRINT ---------- TU ---------- (AND CORRECT SPELLING)
PRINT ---------- X ---------- (AND CORRECT SPELLING)
IS LIST COMPLETE ?
YES
NO
TRANSFER DATA TO LESSON OPTIMIZATION PROGRAM
SYSTEM PREPARES TO PRESENT NEXT WORD
STUDENT SIGN ON ROUTINE AT START OF DAY
STUDENT SIGN OFF ROUTINE

Figure 2: Daily list presentation routine.

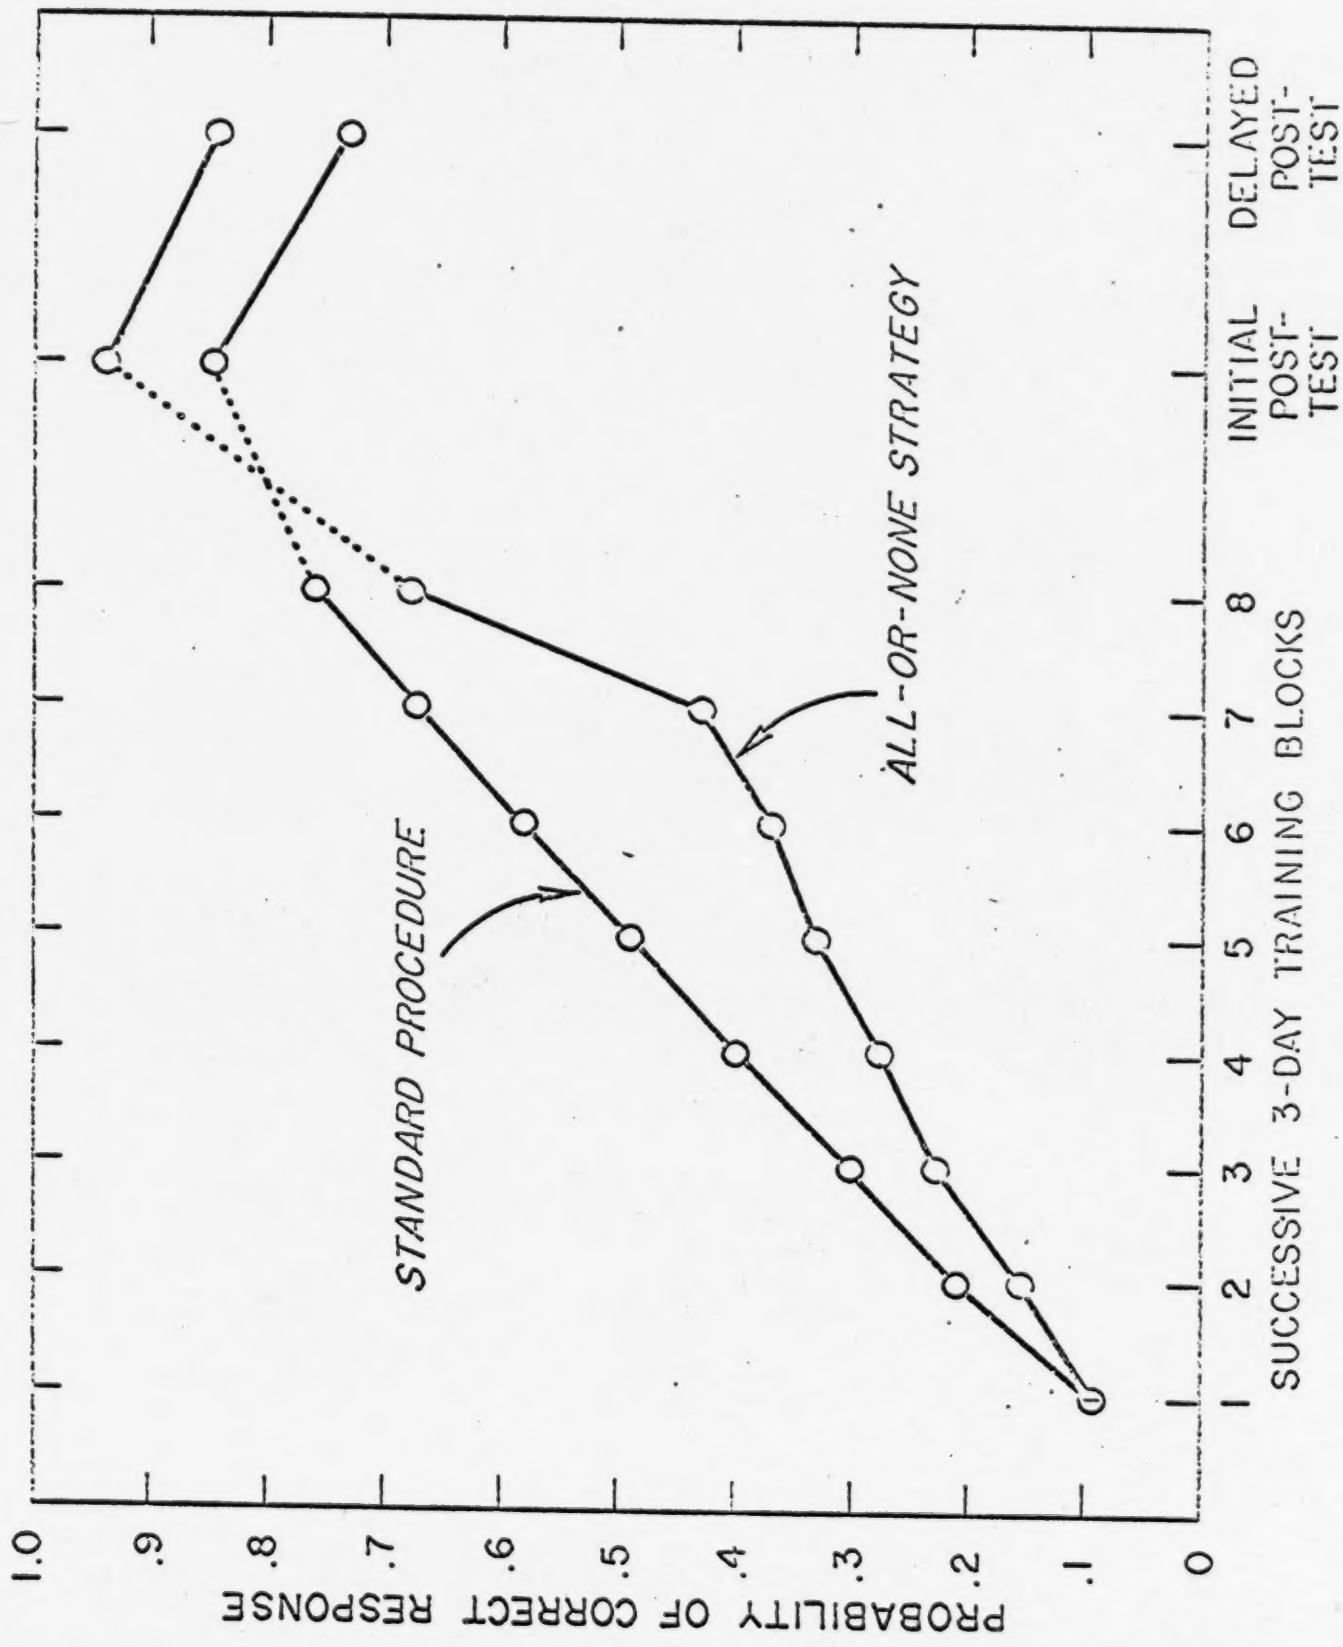

Figure 3: Probability of correct response in Lorton's experiment.

so it is necessary to consider strategies based on more general models. Laubsch (1969), in a Ph.D. Thesis conducted under the auspices of the present grant, considered the optimization problem for cases where the RTI model is appropriate. He made what is perhaps a more significant departure from the assumptions of the all-or-none strategy by allowing the parameters of the model to vary with students and items. The following discussion is based upon Laubsch's work, but introduces a more satisfactory formulation of individual differences. This change and the estimation of initial condition parameters produce experimental measures of the effectiveness of optimization procedures that are significantly greater than those reported by Laubsch.

It is not difficult to derive an approximation to the optimal strategy for the RTI model that can accommodate student and item differences in parameter values, if these parameters are known. Since parameter values must be specified in order to make the necessary calculations to determine the optimal study list, it makes little difference whether these numbers are fixed or vary with students and items. However, making estimates of these parameter values in the heterogeneous case presents some difficulties.

When the parameters of a model are homogeneous, it is possible to pool data from different subjects and items to obtain precise estimates. Estimates based on a sample of students and items can be used to predict the performance of other students or the same students on other items. When the parameters are heterogeneous, these advantages no longer exist unless variations in the parameter values take some known form. For this reason it is necessary to formulate a model stating the composition of each parameter in terms of a subject and item component.

Let π_{ij} be a generic symbol for a parameter characterizing student i and item j. An example of the kind of relationship desired is a fixed-effects subjects-by-items analysis of variance model:

$$E(\pi_{ij}) = m + a_i + d_j \tag{7}$$

Where m is the mean, a_i is the ability of student i, and d_j is the difficulty of item j. Because the learning model parameters we are interested in are probabilities, the above assumption of additivity is not met; that is, there is no guarantee that Equation 7 would yield estimates bounded between 0 and 1. But there is a transformation of the parameter that circumvents this difficulty. In the present context, this transformation has an interesting intuitive justification.

Instead of thinking directly in terms of the parameter π_{ij}, it is helpful to think in terms of the odds ratio $\pi_{ij}/1-\pi_{ij}$. Allow two assumptions: (1) the odds ratio is proportional to student ability; (2) the odds ratio is inversely proportional to item difficulty. This can be expressed algebraically as

$$\frac{\pi_{ij}}{1-\pi_{ij}} = \kappa \frac{a_i}{d_j}, \tag{8}$$

where κ is a proportionality constant. Taking logarithms on both sides yields

$$\log \frac{\pi_{ij}}{1-\pi_{ij}} = \log \kappa + \log a_i - \log d_j \,. \tag{9}$$

The logarithm of the odds ratio is usually referred to as the "logit." Let $\log \kappa = \mu$, $\log a_i = A_i$, and $-\log d_j = D_j$. Then Equation 9 becomes

$$\text{logit } \pi_{ij} = \mu + A_i + D_j \,. \tag{10}$$

Thus, the two assumptions made above lead to an additive model for the values of the parameters transformed by the logit function. Equation 10, by defining a subject-item parameter π_{ij} in terms of a subject parameter A_i applying to all items and an item parameter D_j applying to all subjects, significantly reduces the number of parameters to be estimated. If there are N items and S subjects, then the model requires only N+S parameters to specify the learning parameters for NxS subject-items. More importantly, it makes it possible to predict a student's performance on items he has not been exposed to from the performance of other students on them. This formulation of learning parameters is essentially the same as the treatment of an analogous problem in item analysis given by Rasch (1966). Discussion of this and related models for problems in mental test theory is given by Birnbaum (1968).

Given data from an experiment, Equation 10 can be used to obtain reasonable parameter estimates, even though the parameters vary with students and items. The parameters π_{ij} are first estimated for each student-item protocol, yielding a set of initial estimates. Next the logistic transformation is applied to these initial estimates, and then using these values subject and item effects (A_i and D_j) are estimated by standard analysis of variance procedures. The estimates of student and item effects are used to adjust the estimate of each transformed student-item parameter, which in turn is transformed back to obtain the final estimate of the original student-item parameter.

The first students in an instructional program which employs a parameter-dependent optimization scheme like the one outlined above do not benefit maximally from the program's sensitivity to individual differences in students and items; the reason is that the initial parameter estimates must be based on the data from these students. As more and more students complete the program, estimates of the D_j's become more precise until finally they may be regarded as known constants of the system. When this point has been reached, the only task remaining is to estimate A_i for each new student entering the program. Since the D_j's are known, the estimates of π_{ij} for a new student are of the right order, although they may be systematically high or low until the student component can be accurately assessed.

Parameter-dependent optimization programs with the adaptive character just described are potentially of great importance in long-term instructional programs. Of interest here is the RTI model, but the method of decomposing parameters into student and item components would apply to other models as well. We turn now to an experimental test of the adaptive optimization

program based on the RTI model. In this case the parameters α, c, and q' of the RTI model were separated into item and subject components following the logic of Equation 10. That is, the parameters for subject i working on item j were defined as follows:

$$\text{logit } \alpha_{ij} = \mu^{(\alpha)} + A_i^{(\alpha)} + D_j^{(\alpha)}$$

(11)

$$\text{logit } c_{ij} = \mu^{(c)} + A_i^{(c)} + D_j^{(c)}$$

$$\text{logit } q'_{ij} = \mu^{(q')} + A_i^{(q')} + D_j^{(q')}$$

Note that $A_i^{(\alpha)}$, $A_i^{(c)}$ and $A_i^{(q')}$ are measures of the ability of subject i and hold for all items, whereas $D_j^{(\alpha)}$, and $D_j^{(c)}$ and $D_j^{(q')}$ are measures of the difficulty of item j and hold for all subjects.

The instructional program was designed to teach 300 Swahili vocabulary items to college-level students. Two presentation strategies were employed: (1) the all-or-none procedure and (2) the adaptive optimization procedure based on the RTI model. As in the Lorton study, a within-subjects design was employed in order to provide a sensitive comparison of the strategies. For each student two sublists of 150 items were formed at random from the master list; instruction on items from one sublist was governed by the all-or-none strategy, and by the adaptive optimization strategy for the other sublist. Each day a student was tested on and studied 100 items presented in a random order; 50 items were from the all-or-none sublist chosen using the all-or-none strategy, and 50 from the adaptive optimization list selected according to that strategy. A Swahili word would be displayed and the student was required to give its English translation. Reinforcement consisted of a printout of the correct Swahili-English pair. Twenty such training sessions were involved, each lasting for approximately one hour. Two or three days after the last training session an initial post-test was administered over all 300 items; a delayed post-test was given approximately two weeks later.

The lesson optimization program for the RTI model was more complex than those described earlier. Each night the response data for that day was entered into the system and used to update parameter estimates; in this case an exact record of the complete presentation sequence and response history had to be preserved. A computer-based search algorithm was used to estimate parameters and thus the more accurate the previous day's estimates, the more rapid was the search for the updated parameter values. Once updated estimates had been obtained, they were entered into the optimization program to select individual lists for each student to be run the next day. Early in the experiment (before estimates of $D^{(\alpha)}$, $D^{(c)}$ and $D^{(q')}$ had stabilized) the computation time was fairly lengthy, but it rapidly decreased as more data accumulated and the system homed in on precise estimates of item difficulty.

Figure 4 presents the final test results and indicates that for both the initial and delayed post-tests the parameter-dependent strategy of the RTI model was markedly superior to the all-or-none strategy; on the initial post-test the relative improvement was 41 percent and 67 percent on the delayed post-test. It is apparent that the parameter-dependent strategy was more sensitive than the all-or-none strategy in identifying and presenting those items that would benefit most from additional study. Another feature of the experiment was that students were run in successive groups, each starting about one week after the prior group. As the theory would predict, the overall gains produced by the parameter-dependent strategy increased from one group to the next. The reason is that early in the experiment estimates of item difficulty were crude, but improve with each successive wave of students. Near the end of the experiment estimates of item difficulty were quite exact, and the only task that remained when a new student came on the system was to estimate his particular $A^{(\alpha)}$, $A^{(c)}$, and $A^{(q')}$ values.

Another set of experiments dealing with a similar problem is presented in the appendix to this report. These experiments are particularly important because they examine the issue of learner-controlled instruction as a supplement to strategies of the sort considered above.

CONCLUSIONS AND RECOMMENDATIONS

The studies reported here illustrate one approach that can contribute to the development of a theory of instruction. They deal with relatively simple problems and thus do not indicate the range of developments that are clearly possible. It would be a mistake, however, to conclude that this approach offers a solution to the problems facing education. There are some fundamental obstacles that limit the generality of the work.

The major obstacles may be identified in terms of the four criteria we specified as prerequisites for an optimal strategy. The first criterion concerns the formulation of learning models. The models that now exist are totally inadequate to explain the subtle ways by which the human organism stores, processes, and retrieves information. Until we have a much deeper understanding of learning, the identification of truly effective strategies will not be possible. However, an all-inclusive theory of learning is not a prerequisite for the development of optimal procedures. What is needed instead is a model that captures the essential features of that part of the learning process being tapped by a given instructional task. Even models that may be rejected on the basis of laboratory investigation can be useful in deriving instructional strategies. The two learning models considered in this paper are extremely simple, and yet the optimal strategies they generate are quite effective. My own preference is to formulate as complete a learning model as intuition and data will permit and then use that model to investigate optimal procedures; when possible the learning model will be represented in the form of mathematical equations but otherwise as a set of statements in a computer-simulation program. The main point is that the development of a theory of instruction cannot progress if one holds the view that a complete theory of learning is a prerequisite. Rather, advances in learning theory will affect the development of a theory of instruction, and conversely the development of a theory of instruction will influence research on learning.

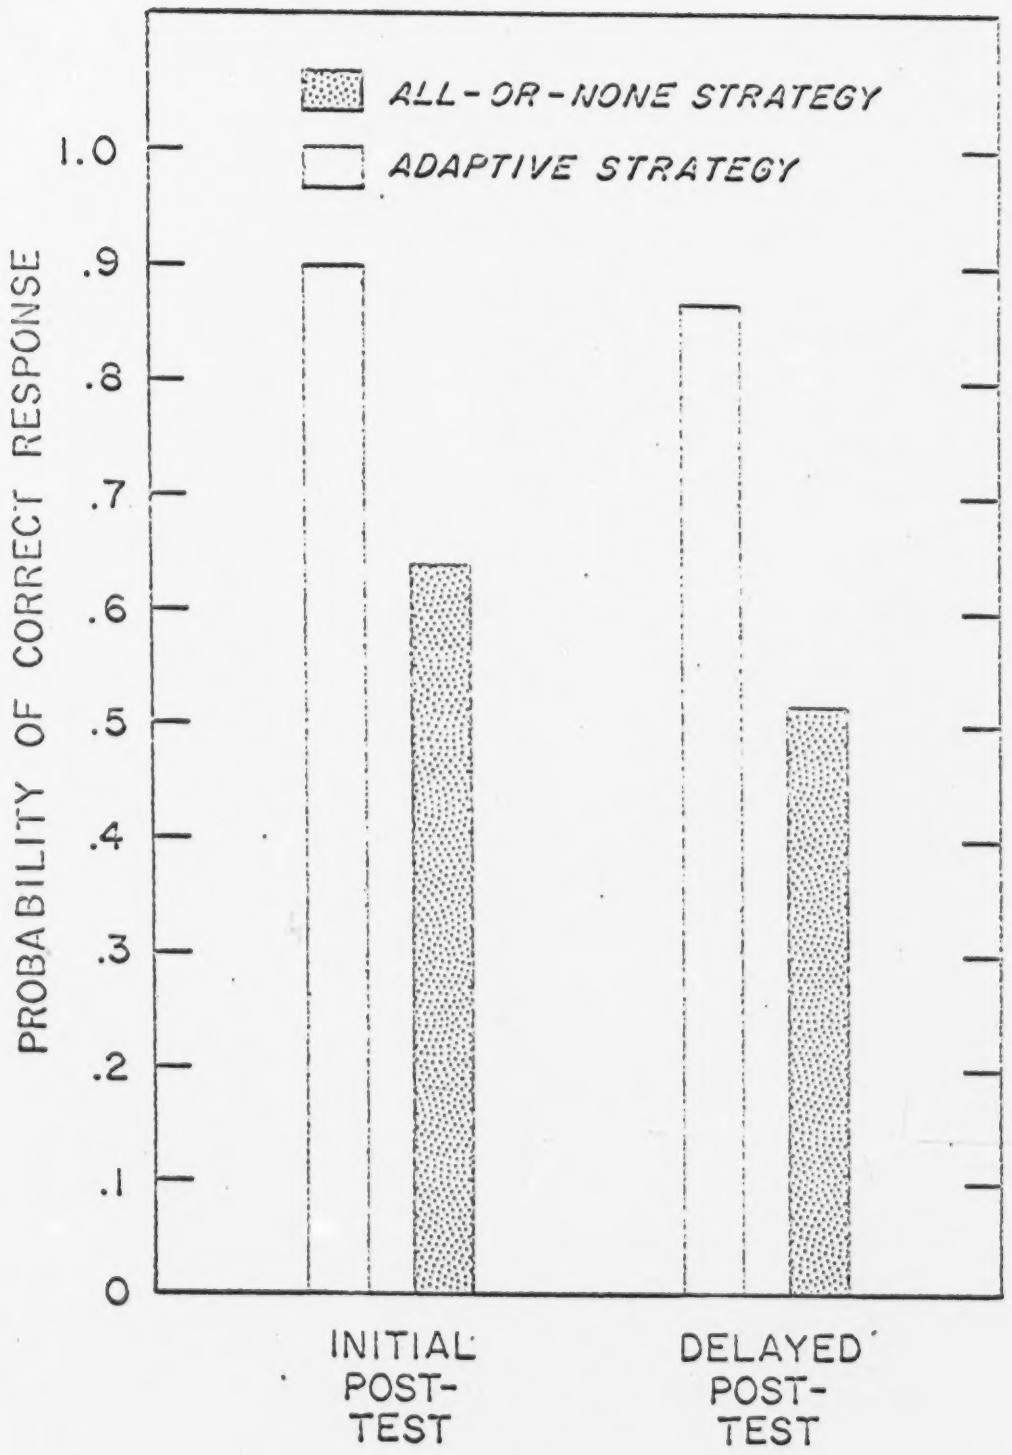

Figure 4: Post-test performance for the all-or-none strategy and for the parameter-dependent strategy of the RTI model.

The second criterion for deriving an optimal strategy requires that admissible instructional actions be clearly specified. The set of potential instructional inputs places a definite limit on the effectiveness of the optimal strategy. In my opinion powerful instructional strategies must necessarily be adaptive; that is, they must be sensitive on a moment-to-moment basis to a learner's unique response history. My judgment on this matter is based on limited experience, restricted primarily to research on teaching initial reading. In this area, however, the evidence seems to be absolutely clear: the manipulation of method variables accounts for only a small percentage of the variance when not accompanied by instructional strategies that permit individualization. Method variables like the modified teaching alphabet, oral reading, the linguistic approach, and others undoubtedly have beneficial effects. However, these effects are minimal in comparison to the impact that is possible when instruction is adaptive to the individual learner. Significant progress in dealing with the nation's problem of teaching reading will require individually prescribed programs, and sophisticated programs will necessitate some degree of computer intervention either in the form of CAI or computer-managed instruction. As a corollary to this point, it is evident from observations of students on our CAI Reading Program that the more effective the adaptive strategy the less important are extrinsic motivators. Motivation is a variable in any form of learning, but when the instructional process is truly adaptive the student's progress is sufficient reward in its own right.

The third criterion for an optimal strategy deals with instructional objectives, and the fourth with cost-benefit measures. In the analyses presented here, it was tacitly assumed that the curriculum material being taught is sufficiently beneficial to justify allocating time to it. Further, in both examples the costs of instruction were assumed to be the same for all strategies. If the costs of instruction are equal for all strategies, they may be ignored and attention focused on the comparative benefits of the strategies. This is an important point because it greatly simplifies the analysis. If both costs and benefits are significant variables, then it is essential that both be accurately estimated. This is often difficult to do. When one of these quantities can be ignored, it suffices if the other can be assessed accurately enough to order the possible outcomes. As a rule, both costs and benefits must be weighed in the analysis, and frequently subtopics within a curriculum vary significantly in their importance. In some cases, whether or not a certain topic should be taught at all is the critical question. Smallwood (1971) has treated problems similar to the ones considered in this article in a way that includes some of these factors in the structure of costs and benefits.

My last remarks deal with the issue of learner-controlled instruction. One way to avoid the challenge and responsibility of developing a theory of instruction is to adopt the view that the learner is the best judge of what to study, when to study, and how to study. I am alarmed by the number of individuals who advocate this position despite a great deal of negative evidence. Don't misinterpret this remark. There obviously is a place for the learner's judgments in making instructional decisions. In several CAI programs that I have helped develop, the learner plays an important role in determining the path to be followed through the curriculum. However, using

the learner's judgment as one of several items of information in making an instructional decision is quite different from proposing that the learner should have complete control. Our data, and the data of others, indicate that the learner is not a particularly effective decision maker. Arguments against learner-controlled programs are unpopular in the present climate of opinion, but they need to be made so that we will not be seduced by the easy answer that a theory of instruction is not required because, "who can be a better judge of what is best for the student than the student himself."

It has become fashionable in recent years to criticize learning theorists for ignoring the prescriptive aspects of instruction, and some have argued that efforts devoted to the laboratory analysis of learning should be redirected to the study of learning as it occurs in real-life situations. These criticisms are not entirely unjustified for in practice psychologists have too narrowly defined the field of learning, but to focus all effort on the study of complex instructional tasks would be a mistake. Some successes might be achieved, but in the long run understanding complex learning situations must depend upon a detailed analysis of the elementary perceptual and cognitive processes that make up the human information handling system. The trend to press for relevance of learning theory is healthy, but if the surge in this direction goes too far, we will end up with a massive set of prescriptive rules and no theory to integrate them.

It needs to be emphasized that the interpretation of complex phenomena is problematical, even in the best of circumstances. The case of hydrodynamics is a good example for it is one of the most highly developed branches of theoretical physics. Differential equations expressing certain basic hydrodynamic relationships were formulated by Euler in the eighteenth century. Special cases of these equations sufficed to account for a wide variety of experimental data. These sucdesses prompted Lagrange to assert that the success would be universal were it not for the difficulty in integrating Euler's equations in particular cases. Lagrange's view is still widely held, in spite of numerous experiments yielding anamolous results. Euler's equations have been integrated in many cases, and the results were found to disagree dramatically with observation, thus contradicting Lagrange's assertion. The problems involve more than mere fine points, and raise serious paradoxes when extrapolations are made from results obtained in the laboratory to actual conditions. The following quotation from Birkhoff (1960) should strike a sympathetic cord among those trying to relate psychology and education: "These paradoxes have been the subject of many witticisms. Thus, it has been said that in the nineteenth century, fluid dynamicists were divided into hydraulic engineers who observed what could not be explained, and mathematicians who explained things that could not be observed. It is my impression that many survivors of both species are still with us."

Research on learning appears to be in a similar state. Educational researchers are concerned with experiments that cannot be readily interpreted in terms of learning theoretic concepts, while psychologists continue to develop theories that seem to be applicable only to the phenomena of the laboratory. Hopefully, work of the sort described here will bridge this gap and help lay the foundations for a theory of instruction.

REFERENCES

Atkinson, R. C. Computer-assisted learning in action. *Proceedings of the National Academy of Sciences*, 1969, *63*, 588-594.

Atkinson, R. C., & Crothers, E. J. A comparison of paired-associate learning models having different acquisition and retention axioms. *Journal of Mathematical Psychology*, 1964, *2*, 285-315.

Atkinson, R. C., Fletcher, J. D., Chetin, H. C., & Stauffer, C. M. *Instruction in initial reading under computer control: The Stanford project*. Educational Technology Research Reports Series, Englewood Cliffs, N.J.: Educational Technology Publications, 1971.

Bellman, R. *Adaptive control processes*. Princeton: Princeton University Press, 1961.

Birkhoff, G. *Hydrodynamics: A study in logic, fact, and similitude*. Second edition. Princeton: Princeton University Press, 1960.

Birnbaum, A. Some latent trait models and their use in inferring an examinee's ability. In F. M. Lord and M. R. Novick (Eds.), *Statistical theories of mental test scores*. Reading, Mass.: Addison-Wesley, 1968. Chapters 17-20.

Bruner, J. S. Some theorems on instruction stated with reference to mathematics. In E. R. Hilgard (Ed.), *Theories of learning and instruction: The Sixty-Third Yearbook of the National Society for the Study of Education*. Chicago: The University of Chicago Press, 1964. Pp. 306-335.

Calfee, R. C. The role of mathematical models in optimizing instruction. *Scientia: Revue Internationale de Synthèse Scientifique*, 1970, *105*, 1-25.

Dear, R. E., Silberman, H. F., Estavan, D. P., & Atkinson, R. C. An optimal strategy for the presentation of paired-associate items. *Behavioral Science*, 1967, *12*, 1-13

Green, G. J., & Atkinson, R. C. Models for optimizing the learning process. Psychological Bulletin, 1966, 66, 309-320.

Hilgard, E. R. (Ed.) Theories of learning and instruction: The Sixty-Third Yearbook of the National Society for the Study of Education. Chicago: The University of Chicago Press, 1964.

Karush, W., & Dear, R. E. Optimal stimulus presentation strategy for a stimulus sampling model of learning. Journal of Mathematical Psychology, 1966, 3, 19-47.

Laubsch, J. H. An adaptive teaching system for optimal item allocation. Unpublished doctoral dissertation, Stanford University, 1969.

Lorton, P., Jr. Computer-based instruction in spelling: An investigation of optimal strategies for presenting instructional material. Unpublished doctoral dissertation. Stanford University, 1971.

Norman, M. F. Incremental learning on random trials. Journal of Mathematical Psychology, 1964, 2, 336-350.

Raiffa, H., & Schlaiffer, R. Applied statistical decision theory. Cambridge, Mass.: MIT Press, 1968.

Rasch, G. An individualistic approach to item analysis. In P. F. Lazarsfeld and N. W. Henry (Eds.), Readings in mathematical social science. Chicago: Science Research Associates, 1966. Pp. 89-107.

Smallwood, R. D. The analysis of economic teaching strategies for a simple learning model. Journal of Mathematical Psychology, 1971, in press.

APPENDIX I: OPTIMIZING THE LEARNING OF A SECOND-LANGUAGE VOCABULARY

Article accepted for publication in the

Journal of Experimental Psychology

Abstract

The problem is to optimize the learning of a large German-English vocabulary. Four optimization strategies are proposed and evaluated experimentally. The first strategy involves presenting items in a random order and serves as a benchmark against which the others can be evaluated. The second strategy permits S to determine on each trial of the experiment which item is to be presented, thus placing instruction under "learner control." The third and fourth strategies are based on a mathematical model of the learning process; these strategies are computer controlled and take account of S's response history in making decisions about which items to present next. Performance on a delayed test administered one week after the instructional session indicated that the learner-controlled strategy yielded a gain of 53% when compared to the random procedure, whereas the best of the two computer-controlled strategies yielded a gain of 108%. Implications of the work for a theory of instruction are considered.

OPTIMIZING THE LEARNING OF A SECOND-LANGUAGE VOCABULARY

Richard C. Atkinson

Stanford University

Stanford, California 94305

This paper examines the problem of individualizing the instructional sequence so that the learning of a second-language vocabulary occurs at a maximum rate. The constraints imposed on the experimental task are those that typically apply to vocabulary learning in an instructional laboratory. A large set of German-English items are to be learned during an instructional session which involves a series of discrete trials. On each trial one of the German words is presented and S attempts to give the English translation; the correct translation is then presented for a brief study period. A predetermined number of trials is allocated for the instructional session, and after some intervening period of time a test is administered over the entire vocabulary set. The problem is to specify a strategy for presenting items during the instructional session so that performance on the delayed test will be maximized. The instructional strategy will be referred to as an adaptive teaching system to the extent that it takes into account S's response history in deciding which items to present from trial to trial.

In this paper four strategies for sequencing the instructional material are considered. One strategy (designated RO) is to cycle through the set of items in a random order; this strategy is not expected to be particularly effective but it provides a benchmark against which to evaluate other procedures. A second strategy (designated SS) is to let S determine for himself

how best to sequence the material. In this mode S decides on each trial which item is to be tested and studied; the learner rather than an external controller determines the sequence of instruction. The third and fourth sequencing schemes (designated OE and OU) can be regarded as adaptive teaching systems and are based on a formal analysis of the learning process. If a mathematical model of the learning process can be stated then it is possible, at least in theory, to derive an optimal strategy. In this paper two instructional strategies derived from a mathematical learning model are examined. The details of these strategies will be presented later.

Before proceeding further, it will be useful to provide an overview of the experimental task. The experiment is run under computer control and involves the learning of a set of 84 German-English items. The Ss are required to participate in two sessions: an instructional session and a test session administered one week later. The delayed test is the same for all Ss and involves a test over the entire set of items. The instructional session is more complicated. The vocabulary items are divided into seven lists each containing twelve German words; the lists are arranged in a round-robin order (see Fig. 1). On each trial of the instructional session a list is displayed and S is permitted to inspect it for a brief period of time. Then one of

---

Insert Figure 1 about here

---

the items on the displayed list is identified for test. In the RO, OE and OU conditions the item is selected by the computer; in the SS condition the item is self-selected by S. After an item has been selected for test, S attempts to provide a translation; then feedback regarding the correct translation is

given. The next trial begins with the computer displaying the next list in the round-robin and the same procedure is repeated. The experiment continues in this fashion for 336 trials (see Fig. 2).

Insert Figure 2 about here

The concern of the experiment is to evaluate the relative effectiveness of the four instructional strategies. Of particular interest is whether strategies derived from a theoretical analysis of the learning process can be as effective as a procedure where S makes his own decisions. If, in fact, the learner is his own best decision maker then the educator's problems are simplified; the appropriate prescription is to place more instruction under learner control.

METHOD

Subjects.- The Ss were 120 undergraduates enrolled in the summer session at Stanford University; 30 Ss were randomly assigned to each of the four experimental conditions. None of the students had prior course work in German and none professed familiarity with the language. The Ss were run in groups of eight with two Ss in each group assigned to one of the four experimental conditions.

Materials.- Seven lists of 12 German words per list were formed. Fig. 1 displays one of the lists as it was presented to S. Based on prior experimentation the lists were judged to be of roughly equal difficulty. All words were concrete nouns typically taught during the first course in German.

Procedure.- The experiment was conducted in the Computer-Based Learning Laboratory at Stanford University. The control functions were performed by programs run on a modified PDP-1 computer manufactured by The Digital Equipment Corp. and under control of a time-sharing system. Eight teletypewriters were housed in a soundproof room and faced a projection screen mounted on the front wall. The instructional session lasted approximately two hours with a 5 min. break in the middle. Each trial was initiated by projecting one of the display lists on the front wall of the room; the list remained on the screen throughout the trial. The Ss were permitted to inspect the list for approximately 10 sec. In the RO, OE and OU conditions this inspection period was followed by the computer typing a number from one to twelve on each S's teletypewriter indicating the item to be tested on that trial; the number typed on a given teletypewriter depended on that S's particular control program. In the SS condition, S typed one of 12 numbered keys during the inspection period to indicate to the computer which item he wanted to be tested on.

At the end of the inspection period S was required to type out the English translation for the designated German word and then strike the "slash" key, or if unable to provide a translation to simply hit the "slash" key. After the "slash" key had been activated the computer typed out the correct translation and spaced down two lines in preparation for the next trial. The trial terminated with the offset of the display list and the next trial began immediately with the onset of the next display list. A complete trial took approximately 20 sec. and the timing of events (within and between trials) was synchroncus for the eight Ss run together. The instructional session involved a total of 336 trials which meant that each list was displayed 48

times. In the RO condition this number of trials permitted each of the items on a list to be tested and studied an average of four times.

The delayed-test session, conducted seven to eight days later, was precisely the same for all Ss. All testing was done on the teletypewriters. A trial began with the computer typing a German word, and S was then required to type the English translation. The 84 German items were presented in a random order and S received no feedback on the correctness of his response. During the delayed-test session the trial sequence was self-paced.

All Ss were told at the beginning of the experiment that there would be a delayed-test session and that their principal goal was to achieve as high a score as possible on that test. They were told, however, not to think about the experiment or rehearse any of the material during the intervening week; these instructions were emphasized at the beginning and end of the instructional session and later reports from Ss confirmed that they made no special effort to rehearse the material during the week between instruction and the delayed test. The instructions emphasized that S should try to provide a translation for every item tested during the instructional session; if S was uncertain but could offer a guess he was encouraged to enter it. In the RO, OE and OU conditions no additional instructions were given. In the SS condition, Ss were told that their trial-to-trial selection of items should be done with the aim of mastering the total list. They were instructed that it was best to test themselves on words they did not know rather than on familiar ones.

RESULTS

The results of the experiment are summarized in Fig. 3. On the left side of the figure data are presented for performance during the instructional session; on the right side are results from the delayed test. The data from the instructional session are presented in four successive blocks of 84 trials

---
Insert Figure 3 about here
---

each; for the RO condition this means that on the average each item was presented once in each of these blocks. Note that performance during the instructional session is best for the RO condition, next best for the OE condition which is slightly better than the SS condition, and poorest for the OU condition; these differences are highly significant, $F(3,116) = 21.3$, $p < .001$. The order of the experimental groups on the delayed test is completely reversed. The OU condition is by far best with a correct response probability of .79; the SS condition is next with .58 followed closely by the OE condition at .54; the RO condition is poorest at .38 ($F(3,116) = 18.4$, $p < .001$). The observed pattern of results is what one would expect. In the SS condition Ss are trying to test themselves on items they do not know; consequently, during the instructional session, they should have a lower proportion of correct responses than Ss run on the RO procedure where items are tested at random. Similarly, the OE and OU conditions involve a procedure that attempts to identify and test those items that have not yet been mastered and also should produce high error rates during the instructional session. The ordering of groups on the delayed test is reversed since all words are tested in a non-selective fashion; under these conditions a "true" measure of S's mastery of the list is obtained.

The magnitude of the effects observed on the delayed test are large and of practical significance. The SS condition (when compared to the RO condition) leads to a relative gain of 53%, whereas the OU condition yields a relative gain of 108%. It is interesting that S can be very effective in determining an optimal study sequence, but not as effective as the best of the two adaptive teaching systems.

DISCUSSION

At this point we turn to an account of the theory on which the OU and OF schemes are based. Both schemes assume that acquisition of a second-language vocabulary can be described by a fairly simple learning model. It is postulated that a given item is in one of three states (P, T and U) at any moment in time. If the item is in state P then its translation is known and this knowledge is "relatively" permanent in the sense that the learning of other items will not interfere with it. If the item is in state T then it is also known but on a "temporary" basis; in state T the learning of other items can give rise to interference effects that cause the item to be forgotten. In state U the item is not known and S is unable to give a translation. Thus in states P and T a correct translation is given with probability one, whereas in state U the probability is zero.

When item i is presented for test and study the following transition matrix describes the possible change in state from the onset of the trial to its termination:

$$\underline{A}_i = \begin{array}{c} \\ P \\ T \\ U \end{array} \begin{array}{c} \begin{array}{ccc} P & T & U \end{array} \\ \begin{bmatrix} 1 & 0 & 0 \\ x_i & 1-x_i & 0 \\ y_i & z_i & 1-y_i-z_i \end{bmatrix} \end{array}$$

Rows of the matrix represent the state of item i at the start of the trial and columns the state at the end of the trial. On a trial when some item other than item i is presented for test and study (whether that item is a member of item i's display list or some other display list) transitions in the learning state of item i also may take place. Such transitions can occur only if S makes an error on the trial; in that case the transition matrix applied to item i is as follows:

$$\underline{F}_i = \begin{array}{c} \\ P \\ T \\ U \end{array} \begin{array}{c} \begin{array}{ccc} P & T & U \end{array} \\ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1-f_i & f_i \\ 0 & 0 & 1 \end{bmatrix} \end{array}$$

Basically, the idea is that when some other item is presented to which S makes an error (i.e., an item in state U) then forgetting may occur for item i if it is in state T.

To summarize, when item i is presented for test and study transition matrix $\underline{A}_i$ is applied; when some other item is presented that elicits an error then matrix $\underline{F}_i$ is applied. The above assumptions provide a complete account of the learning process. For the task considered in this paper it is also assumed that item i is either in state P (with probability g_i) or in state U (with probability $1-g_i$) at the start of the instructional session; S

either knows the correct translation without having studied the item or does not. The parameter vector $\phi_i = [x_i, y_i, z_i, f_i, g_i]$ characterizes the learning of a given item i in the vocabulary set. The first three parameters characterize the acquisition process; the next parameter, forgetting; and the last, S's knowledge prior to entering the experiment.

For a more detailed account of the model the reader is referred to Atkinson and Crothers (1964) and Calfee and Atkinson (1965). It has been shown in a series of experiments that the model provides a fairly good account of vocabulary learning and for this reason it was used to develop an optimal procedure for controlling instruction. We now turn to a discussion of how OE and OU procedures were derived from the model. Prior to conducting the experiment reported in this paper, a pilot study was run using the same word lists and the RO procedure described above. Data from the pilot study were employed to estimate the parameters of the model; the estimates were obtained using the minimum chi square procedures discussed in Atkinson and Crothers (1964). Two separate estimates of parameters were made. In one case it was assumed that the items were equally difficult and data from all 84 items were lumped together to obtain a single estimate of the parameter vector ϕ; this estimation procedure will be called the equal parameter case (E-case), since all items are assumed to be of equal difficulty. In the second case data were separated by items and an estimate of ϕ_i was made for each of the 84 items (i.e., $84 \times 5 = 420$ parameters were estimated); this procedure will be called the unequal parameter case (U-case). In both the U and E cases it was assumed that there were no differences among Ss; this homogeneity assumption regarding learners will be commented upon later. The two sets of parameter estimates

were used to generate the optimization schemes previously referred to as the OE and OU procedures; the former based on estimates from case E and the latter from case U.

In order to formulate an instructional strategy it is necessary to be precise about the quantity to be maximized. For the present experiment the goal is to maximize the total number of items S correctly translates on the delayed test.[1] To do this, we need to specify the theoretical relationship between the state of learning at the end of the instructional session and performance on the delayed test. The assumption made here is that only those items in state P at the end of the instructional session will be translated correctly on the delayed test; an item in state T at the end of the instructional session is presumed to be forgotten during the intervening week. Thus, the problem of maximizing delayed-test performance involves, at least in theory, maximizing the number of items in state P at the termination of the instructional session.

Having numerical values for parameters and knowing S's response history, it is possible to estimate his current state of learning.[2] Stated more precisely, the learning model can be used to derive equations and, in turn, compute the probabilities of being in states P, T and U for each item at the start of trial n, conditionalized on S's response history up to and including trial n-1. Given numerical estimates of these probabilities a strategy for optimizing performance is to select that item for presentation (from the current display list) that has the greatest probability of moving into state P if it is tested and studied on the trial.[3]

The optimization procedure described above was implemented on the computer and permitted decisions to be made on-line for each S on a trial-by-trial basis. For Ss in the OE group, the computations were carried out using the five parameter values estimated under the assumption of homogeneous items (E-case); for Ss in the OU group the computations were based on the 420 parameter values estimated under the assumption of heterogeneous items (U-case).

The OU procedure is sensitive to inter-item differences and consequently generates a more effective optimization strategy than the OE procedure. The OE procedure, however, is not to be ignored for it is nearly as effective as having S make his own instructional decisions, and far superior to a random presentation scheme. If individual differences among Ss also are taken into account, then further improvements in delayed-test performance should be possible; this issue and methods for dealing with individual differences are discussed in Atkinson and Paulson (1972).

The study reported here illustrates one approach that can contribute to the development of a theory of instruction (Hilgard, 1964). This is not to suggest that the OU procedure represents a final solution to the problem of optimal item selection. The model upon which this strategy is based ignores several important factors, such as inter-item relationships, motivation, and short-term memory effects (Atkinson & Shiffrin, 1968, p. 190). Undoubtedly, strategies based on learning models that take these variables into account would yield superior procedures.

Although the task considered in this paper deals with a limited form of instruction, there are at least two practical reasons for studying it. First, this type of task occurs in numerous learning situations; no matter what the

pedagogical orientation, any initial reading program or foreign-language course involves some form of list learning. In this regard it should be noted that a modified version of the OU strategy has been used successfully in the Stanford computer-assisted instruction program in initial reading (Atkinson & Fletcher, 1972). Secondly, the study of such relatively simple tasks that can be understood in detail provides prototypes for analyzing more complex optimization problems. At present, analyses comparable to those reported here cannot be made for many instructional procedures of central interest to educators, but examples of this sort help to clarify the steps involved in devising and testing optimal strategies. For a review of work on optimizing learning and references to the literature see Atkinson and Paulson (1972).

REFERENCES

Atkinson, R. C., & Crothers, E. J. A comparison of paired-associate learning models having different acquisitions and retention axioms. Journal of Mathematical Psychology, 1964, 2, 285-315.

Atkinson, R. C., & Fletcher, J. D. Teaching children to read with a computer. The Reading Teacher, 1972, 25, 319-327.

Atkinson, R. C., & Paulson, J. A. An approach to the psychology of instruction. Psychological Bulletin, 1972, in press.

Atkinson, R. C., & Shiffrin, R. M. Human memory: A proposed system and its control processes. In K. W. Spence and J. T. Spence (Eds.), The Psychology of Learning and Motivation: Advances in Research and Theory, Vol. 2. New York: Academic Press, 1968, 89-155.

Calfee, R. C. The role of mathematical models in optimizing instruction. Scientia: Revue Internationale de Synthèse Scientifique, 1970, 105, 1-25.

Calfee, R. C., & Atkinson, R. C. Paired-associate models and the effect of list length. Journal of Mathematical Psychology, 1965, 2, 255-267.

Dear, R. E., Silberman, H. F., Estavan, D. P., & Atkinson, R. C. An optimal strategy for the presentation of paired-associate items. Behavioral Science, 1967, 12, 1-13.

Hilgard, E. R. (Ed.). Theories of learning and instruction: The Sixty-Third Yearbook of the National Society for the Study of Education. Chicago: The University of Chicago Press, 1964.

Groen, G. J., & Atkinson, R. C. Models for optimizing the learning process. Psychological Bulletin, 1966, 66, 309-320.

Laubsch, J. H. Optimal item allocation in computer-assisted instruction. IAG Journal, 1970, 3, 295-311.

Matheson, J. Optimum teaching procedures derived from mathematical learning models. Report CC52, Institute in Engineering Economic Systems, Stanford University, 1964.

Smallwood, R. C. The analysis of economic teaching strategies for a simple learning model. Journal of Mathematical Psychology, 1971, 8, 285-301.

FOOTNOTES

[1]Other measures can be used to assess the benefits of an instructional strategy; e.g., in this case weights could be assigned to items measuring their relative importance. Also costs may be associated with the various actions taken during an instructional session. Thus, for the general case, the optimization problem involves assessing costs and benefits and finding a strategy that maximizes an appropriate function defined on them. For a discussion of this issue see Atkinson and Paulson (1972), Dear, et al. (1967), and Smallwood (1971).

[2]The S's response history is a record for each trial of the vocabulary item presented and the response that occurred. It can be shown that there exists a sufficient history that contains only the information necessary to estimate S's current state of learning; the sufficient history is always a function of the complete history and the assumed learning model. For the model considered in this paper the sufficient history is fairly simple, but cannot be easily described without extensive notation.

[3]An optimal procedure maximizes the number of items in state P after all trials of the instructional session have been presented. The procedure used here is only a one-stage optimization procedure and there is no guarantee that it is in fact optimal. However, the computations for the N-stage procedure are too time-consuming even for a large computer. Furthermore, a series of Monte Carlo studies indicate that the one-stage procedure is a good approximation to the optimal strategy for a variety of Markov learning models (Matheson, 1964; Laubsch, 1970; Calfee, 1970).

FIGURE CAPTIONS

Figure 1: Schematic representation of the round-robin of display lists and an example of one such list.

Figure 2: Flow chart describing the trial sequence during the instructional session. The selection of a word for test on a given trial (box with heavy border) varied over experimental conditions.

Figure 3: Proportion of correct responses in successive trial blocks during the instructional session and on the delayed test administered one week later.

Signed

Round-robin of Seven Lists

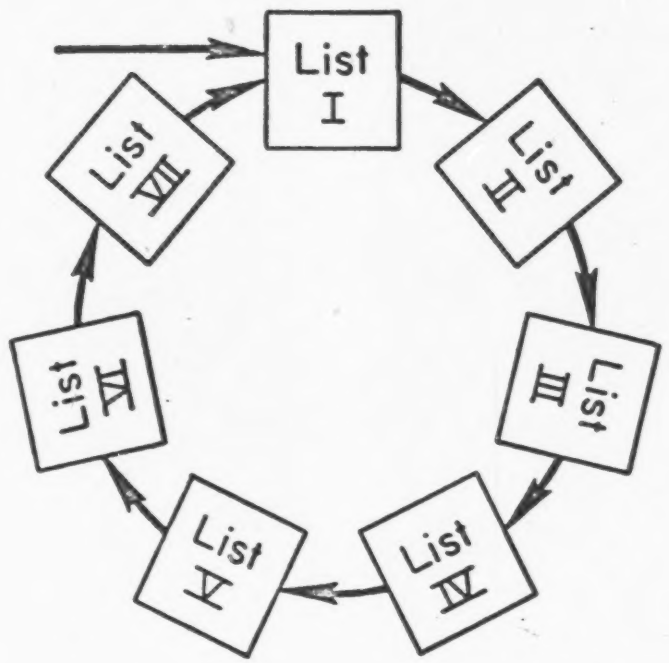

Typical List

1. das Rad
2. die Seite
3. das Kino
4. die Gans
5. der Fluss
6. die Gegend
7. die Kamera
8. der Anzug
9. das Geld
10. der Gipfel
11. das Bein
12. die Ecke

Figure 2

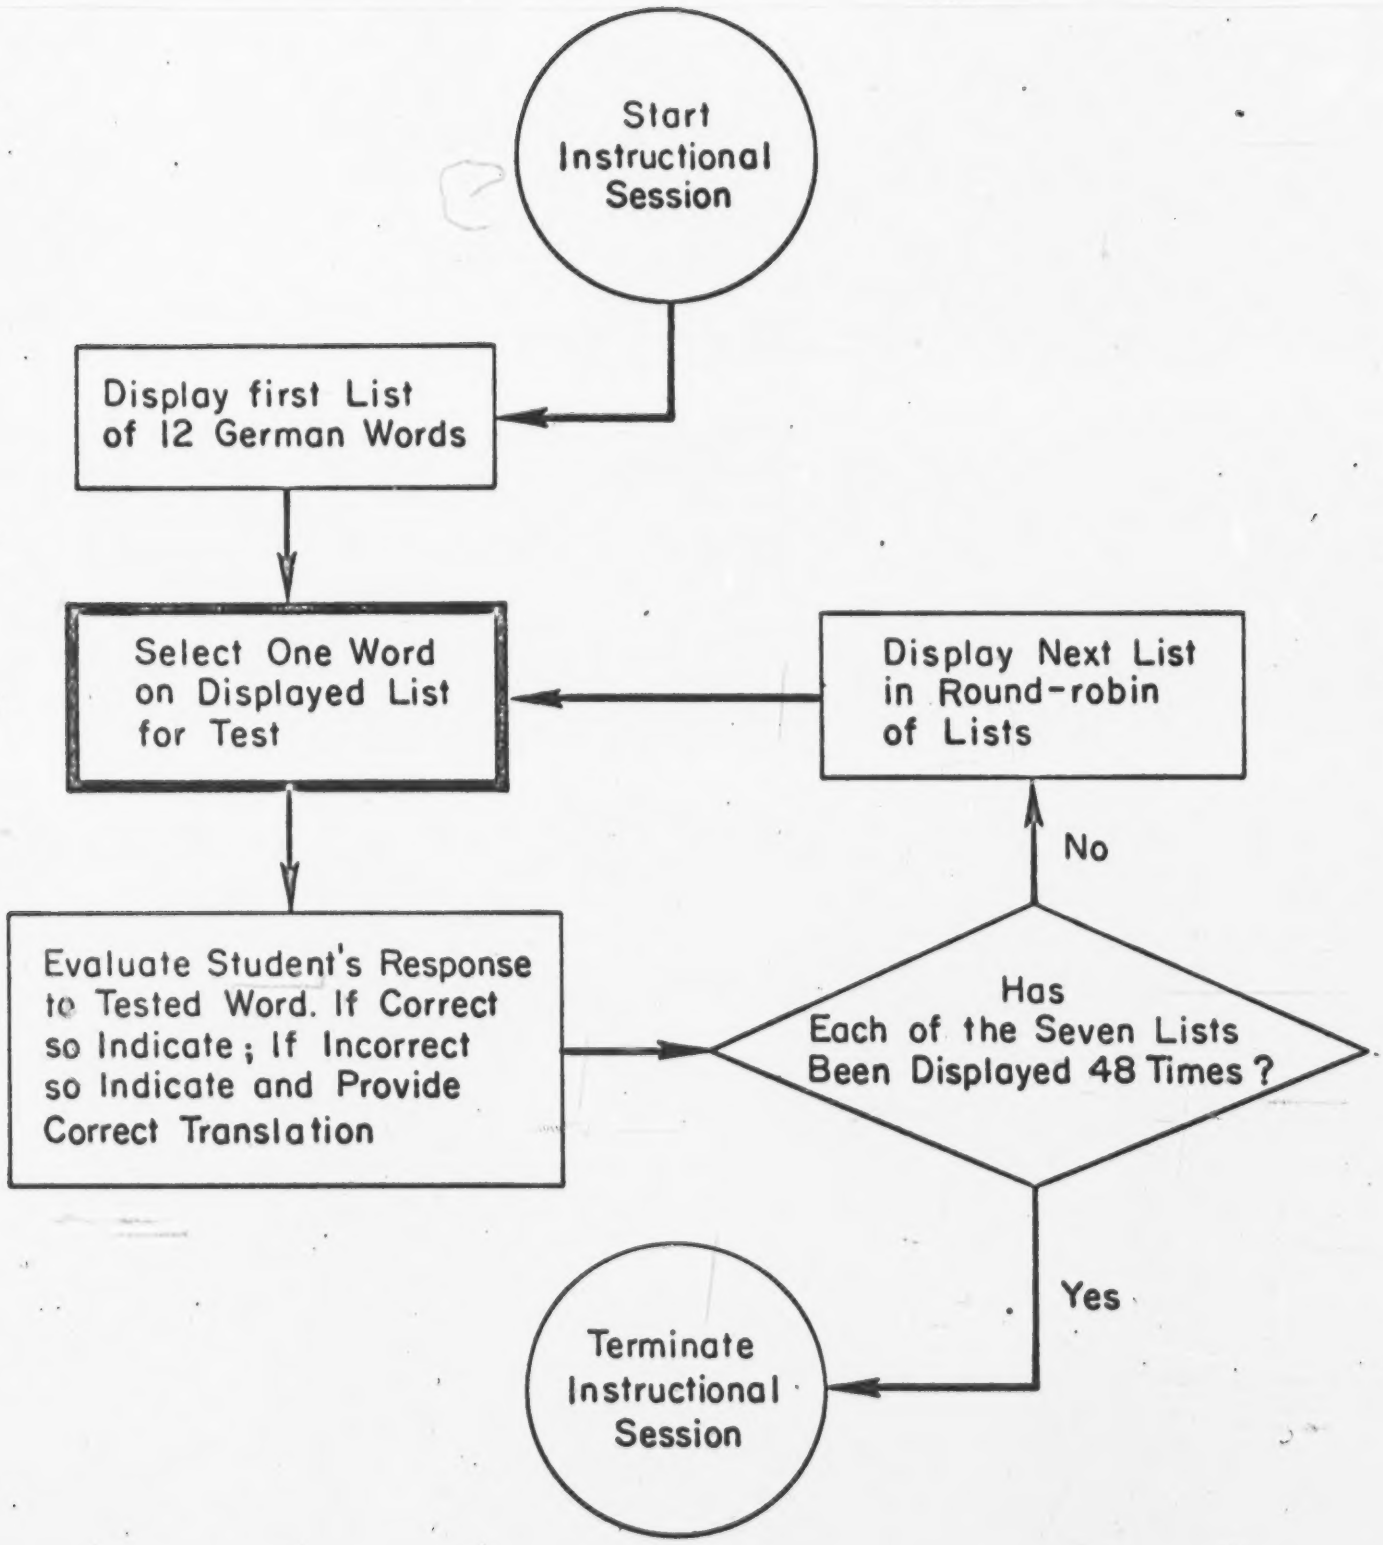
Start Instructional Session
Display first List of 12 German Words
Select One Word on Displayed List for Test
Display Next List in Round-robin of Lists
No
Evaluate Student's Response to Tested Word. If Correct so Indicate; If Incorrect so Indicate and Provide Correct Translation
Has Each of the Seven Lists Been Displayed 48 Times?
Yes
Terminate Instructional Session

Figure 3

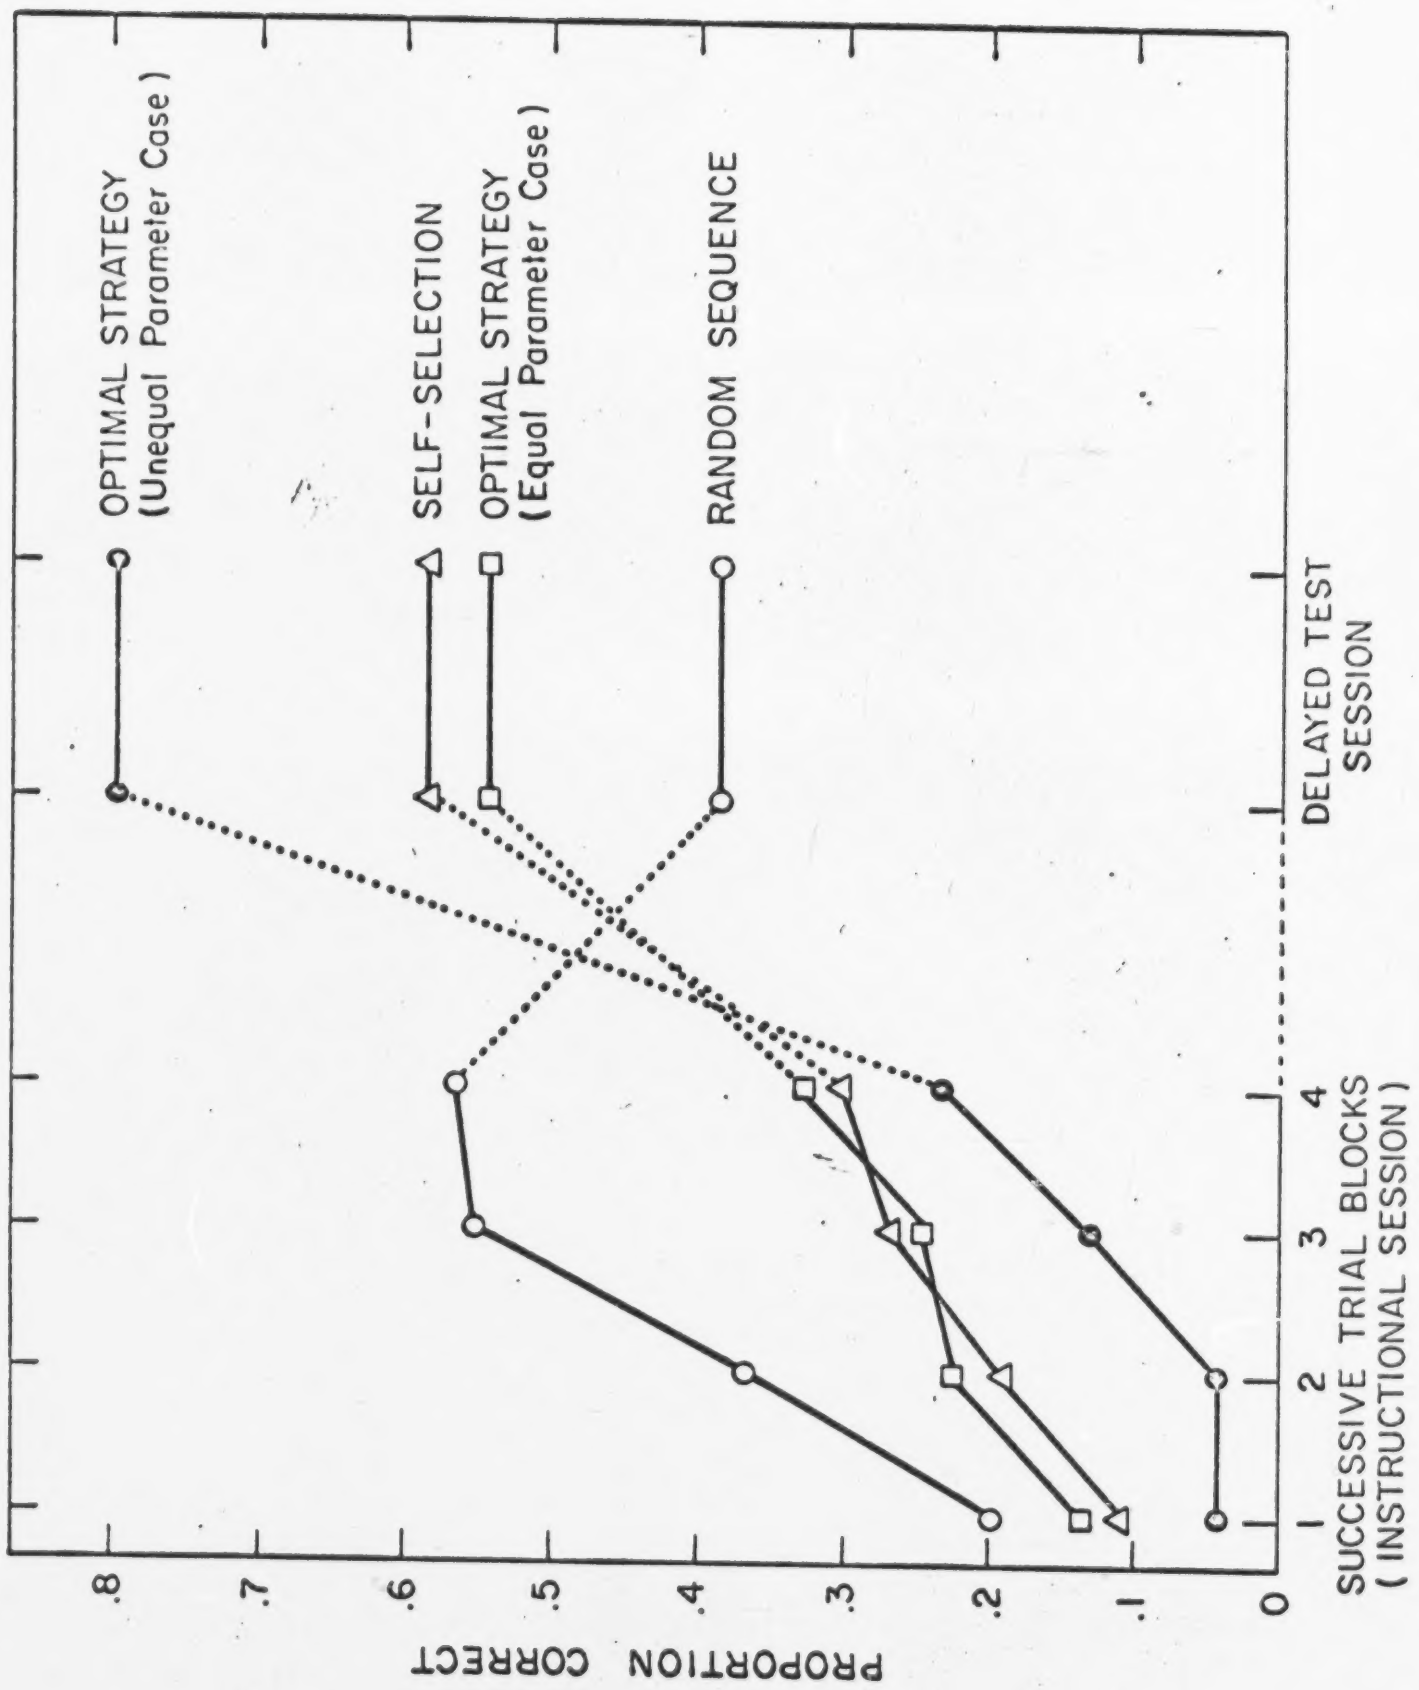
OPTIMAL STRATEGY
(Unequal Parameter Case)
SELF-SELECTION
OPTIMAL STRATEGY
(Equal Parameter Case)
RANDOM SEQUENCE
DELAYED TEST
SESSION
PROPORTION CORRECT
.8
.7
.6
.5
.4
.3
.2
.1
0
1
2
3
4
SUCCESSIVE TRIAL BLOCKS
(INSTRUCTIONAL SESSION)